

Who Captures a Tax Cut? Property Price Capitalization and Fiscal Substitution from France’s Residence Tax Abolition

APEP Autonomous Research*

@olafdrw

March 3, 2026

Abstract

France abolished its €22 billion residence tax (*taxe d’habitation*) between 2018 and 2023. I exploit cross-commune variation in pre-reform tax rates, combining 5.4 million property transactions with fiscal data for 33,000 communes. Despite the reform’s scale, I find no significant cross-sectional capitalization into property prices during 2020–2025: estimates are economically small and statistically indistinguishable from zero. The explanation lies in fiscal substitution. Communes raised property tax (*taxe foncière*) rates by an average of 22 percentage points between 2017 and 2024, largely through mechanical absorption of the département tax share. This fiscal offset explains the absence of net capitalization: homebuyers gained from TH abolition but lost from TFB increases, leaving property prices approximately unchanged. The incidence of tax reform depends critically on the fiscal responses of overlapping governments.

JEL Codes: H22, H71, R31, H77

Keywords: tax capitalization, property tax, fiscal federalism, housing prices, taxe d’habitation, France

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

1. Introduction

Between 2018 and 2023, the French government abolished a €22 billion tax. The *taxe d’habitation*—paid annually by every household occupying a primary residence—was the largest local tax in the country. Standard economic theory makes a sharp prediction about such a reform: property prices should rise by the present value of future tax savings, transferring the windfall to whoever owns the home when the cut is announced (??). Full capitalization turns a recurring tax cut into a one-time wealth transfer. In France, property prices did not respond. This paper explains why.

France recently conducted one of the largest natural experiments on this question. Between 2018 and 2023, the government progressively eliminated the *taxe d’habitation* (TH)—a residence tax paid annually by occupants of their primary dwelling—for all households. The reform was worth approximately €22 billion per year, affecting every primary residence in the country. It was phased by income: 80 percent of households received full exemption by 2020, with the remaining 20 percent phased out through 2023. The reform created massive variation in treatment intensity across communes, since TH rates were locally determined and ranged from near zero to well above 100 percent of the cadastral rental value.

But the French reform has a twist that pure capitalization theory ignores. Communes did not simply absorb the revenue loss. While the central government compensated them with a transfer of national *taxe foncière* (TF) revenue, communes retained the authority to set their own TF rates—and many responded by raising them. Average TF rates increased substantially after 2018, with some communes hiking rates by 10 to 30 percentage points. This fiscal substitution potentially offsets—or even reverses—the benefit of TH abolition for property owners.

This paper provides the first causal analysis of how property markets and local governments jointly respond to a large-scale tax abolition. I construct a continuous-treatment difference-in-differences design that exploits cross-commune variation in pre-reform TH rates as a measure of treatment intensity. Using the universe of French property transactions from the DVF (*Demandes de Valeurs Foncières*) registry—covering approximately one million residential sales per year across more than 30,000 communes—I trace the capitalization dynamics year by year and decompose the net price effect into two components: the gross benefit from TH elimination and the offsetting cost from TF increases.

I find three main results. First, despite the massive scale of the TH abolition, cross-sectional estimates of property price capitalization are economically small and statistically indistinguishable from zero. A one-percentage-point higher pre-reform TH rate—implying a proportionally larger tax cut—is associated with less than a 0.1 percent difference in

property prices per square meter, with confidence intervals that comfortably include zero across all specifications. Second, and more strikingly, communes responded to lost TH revenue with aggressive fiscal substitution: average *taxe foncière* rates rose by approximately 22 percentage points between 2017 and 2024, with the increase strongly predicted by pre-reform TH dependence. This substitution rate of roughly 50 percent implies that about half the TH revenue loss was recouped through higher property taxes on owners. Third, the net fiscal benefit—TH savings minus TFB increases—is correspondingly attenuated, explaining the absence of net capitalization. When I instrument property prices with the net fiscal benefit rather than gross TH savings, the coefficient remains insignificant. The reform transferred fiscal burden from occupants to owners without generating a measurable windfall in property values.

These results challenge three established views. Most directly, this paper advances the long tradition of studying property tax capitalization, from ? through ? and ?. The existing literature has exploited cross-sectional variation in tax rates, changes in assessment practices, and Proposition 13–style rate caps. France’s TH abolition offers an unusually clean setting: the reform was nationally mandated (eliminating concerns about endogenous local tax changes), income-staggered (creating variation in treatment timing), and massive in scale (generating large first-order effects). Moreover, the DVF registry provides universe-level transaction data, enabling precise estimation even in small communes.

Second, this paper contributes to the literature on fiscal federalism and vertical tax competition (??). A standard prediction is that lower-tier governments respond to lost revenue by raising alternative tax instruments—what I call fiscal substitution. While this response is well-documented in federal systems like the United States and Germany (??), it has not been causally estimated in the context of a full tax abolition. The French setting is ideal because the TH was eliminated entirely (not merely reduced), the compensation mechanism was transparent and formula-based, and the TF rate-setting authority was explicitly preserved.

Third, this paper speaks to a growing body of work on French local public finance (???). The only prior study of the TH reform is ?, a policy report for the Institut des Politiques Publiques that uses continuous treatment intensity and finds evidence of price capitalization. My paper differs in three ways: I use the complete DVF universe with over 5 million transactions, I formally estimate the fiscal substitution channel as the primary mechanism rather than a secondary analysis, and I show that the null capitalization result is the *predicted* outcome once fiscal substitution is accounted for—a finding that reconciles the apparent anomaly of a massive tax cut with no price response.

Finally, this paper has implications for tax reform design worldwide. The finding that

fiscal substitution can substantially offset the intended benefits of tax cuts suggests that reformers must consider the full fiscal architecture, not just the abolished tax in isolation. When local governments retain rate-setting authority over complementary tax instruments, the effective incidence of a tax cut depends on the endogenous fiscal response—a mechanism that standard incidence analysis often ignores.

Section 2 provides institutional background. Section 3 presents the conceptual framework. Sections 4–5 describe the data and empirical strategy. Section 6 presents the results. Section 7 discusses mechanisms and implications. Section 8 concludes.

2. Institutional Background

2.1 The Taxe d’Habitation

The *taxe d’habitation* was a local tax levied annually on occupants of residential dwellings in France. Unlike the *taxe foncière* (a property tax paid by owners), the TH was paid by whoever occupied the dwelling on January 1st of each year—whether owner or tenant. The tax base was the cadastral rental value (*valeur locative cadastrale*) of the property, and the rate was set independently by each commune and *intercommunalité* (grouping of communes).

The institutional design of the TH had several distinctive features. First, the tax base—cadastral rental values—had not been comprehensively revised since 1970. While partial adjustments were applied through annual revaluation coefficients, substantial misalignment had accumulated between assessed values and market rents. This created a system in which two otherwise identical properties in neighboring communes could face dramatically different effective tax rates, depending on local rate-setting histories and the idiosyncrasies of the 1970 valuation.

Second, the TH was composed of multiple layers: the commune component, the intercommunal component (*EPCI*), and in some cases a departmental component. Each layer set its own rate independently, so the total TH burden on a household reflected the cumulative decisions of multiple governing bodies. In the analysis, I focus on the commune-level voted rate as the primary measure of local tax intensity, since the commune was the entity with the strongest rate-setting authority and the largest fiscal stake.

Third, the TH applied to both owner-occupants and renters—a critical distinction for the capitalization analysis. When the tax was levied on the occupant, its abolition benefited both groups directly. However, because property prices are determined in the owner market, capitalization theory predicts that the price response depends on how the tax change affects the net return to ownership, not the direct benefit to tenants.

Commune-level voted TH rates varied substantially. In 2017—the last year before the

reform—these rates ranged from near zero to above 150 percent of the cadastral rental value, with a mean of approximately 46 percent and a standard deviation of 29 percentage points. While rates exceeding 100 percent of the cadastral base may seem anomalous, they are a consequence of the 1970 valuation: since the base is far below market rents, even high nominal rates produce modest effective tax burdens. The wide variation reflected decades of autonomous local fiscal decisions and was largely uncorrelated with current economic conditions. The TH represented a major revenue source for communes, generating approximately €22 billion annually (commune and intercommunal layers combined) and accounting for roughly 35 percent of commune tax revenue in the median commune.

2.2 The Abolition Reform (2018–2023)

The abolition of the TH on primary residences was a central campaign promise of Emmanuel Macron during the 2017 presidential election. The reform was implemented progressively through the Finance Laws of 2018–2020:

- **2018:** Households below income thresholds (approximately 80% of the population) received a 30% reduction in their TH bill.
- **2019:** The same households received a 65% reduction.
- **2020:** Full exemption for the initial 80% of households.
- **2021:** The remaining 20% of households received a 30% reduction.
- **2022:** 65% reduction for the remaining households.
- **2023:** Complete abolition of TH on all primary residences.

Several features of the reform are critical for identification. First, the reform was legislated at the national level and implemented uniformly across all communes. Local governments had no discretion over the timing or extent of TH reductions—these were determined entirely by household income and the national schedule. This eliminates the standard endogeneity concern in tax capitalization studies, where local tax changes may be correlated with local economic conditions.

Second, the TH on *secondary residences* (*résidences secondaires*) was *not* abolished. It was maintained at existing rates and, in designated high-demand housing markets (*zones tendues*), communes were authorized to apply surcharges of up to 60%. This creates a natural placebo group: secondary-residence-heavy communes experienced the same local economic conditions but did not receive the TH benefit on a large share of their housing stock.

Third, the income-based staggering of the reform created variation in treatment timing at the commune level. Communes with a higher share of low-income households were effectively treated earlier (their residents received full exemption by 2020), while communes with wealthier populations were treated later (full exemption only in 2023). I construct a “reform share” variable that captures the approximate fraction of households fully exempt in each year: 80 percent in 2020, 90 percent in 2021, 95 percent in 2022, and 100 percent from 2023 onward. This time variation supports the panel event study design.

Fourth, the magnitude of the reform was exceptionally large by international standards. At €22 billion annually, the TH abolition represented the single largest local tax cut in modern French history—comparable in scale to the U.S. Tax Cuts and Jobs Act of 2017 relative to GDP. This scale should have generated detectable capitalization effects even under conservative assumptions about discount rates and housing supply elasticity, making the null result all the more striking.

2.3 Fiscal Compensation and the Tax Foncière Channel

To compensate communes for lost TH revenue, the central government designed a multi-step fiscal transfer mechanism. The key element was the reallocation of the *taxe foncière sur les propriétés bâties* (TFPB) share previously collected by *départements*. Before the reform, each property owner paid TFPB to both the commune and the département (each setting its own rate). Under the compensation scheme, the département TFPB share was transferred to communes, effectively increasing the commune TFPB rate by the former département rate—a *mechanical* component of the observed TFB increase.

The transfer was calibrated to 2017 TH revenue, ensuring that each commune received a fixed nominal amount regardless of subsequent changes in local housing markets or fiscal behavior. The formula was intentionally backward-looking to prevent moral hazard: communes that had maintained low TH rates received less compensation, while those with high TH rates received more.

Crucially, communes retained full authority to set their own TFPB voted rates (*taux votés*) beyond the mechanical transfer. This created an explicit channel for *behavioral* fiscal substitution: communes that lost proportionally more TH revenue could further compensate by raising their own TFPB voted rates above the mechanically adjusted level. The TFPB is paid by property owners (not occupants), so any increase represents a direct cost to homeowners—potentially offsetting the benefit of TH abolition.

The distinction between mechanical and behavioral substitution matters for policy interpretation but not for the capitalization analysis. Whether TFB rates rose because of the département share transfer or because communes voluntarily increased their rates, the

effect on property owners is identical: higher annual TFB bills. A rational buyer pricing a property in 2022 considers the *total* TFB rate, not whether each component was mechanical or discretionary.

Evidence suggests that many communes exercised their discretionary authority. According to the *Cour des comptes*, average TFPB rates increased significantly after 2018, though with substantial heterogeneity across communes (?). The most dramatic example was Paris, which increased its TFPB rate by 52% in a single year (2023), generating widespread media attention and political controversy. However, other communes maintained stable voted rates, creating variation in the degree of behavioral fiscal substitution. This heterogeneity—above and beyond the uniform mechanical transfer—is what drives the cross-sectional correlation between baseline TH rates and TFB changes that I estimate in Section ??.

2.4 The Alsace-Moselle Exception

Communes in the departments of Bas-Rhin (67), Haut-Rhin (68), and Moselle (57) operate under a distinct land registry system (*livre foncier*) inherited from German administration. DVF transaction data is not available for these departments, so they are excluded from the analysis. This affects approximately 5% of the French population.

3. Conceptual Framework

I present a simple framework to organize the empirical analysis. Consider a homebuyer choosing between communes that differ in their TH rates. In a standard Tiebout model with capitalization (??), the purchase price P_c in commune c satisfies a no-arbitrage condition:

$$P_c = V(g_c) - \text{PV}(\tau_c^{TH} + \tau_c^{TF}) \tag{1}$$

where $V(g_c)$ is the value of local public goods, τ_c^{TH} is the annual TH liability, τ_c^{TF} is the annual TF liability, and PV denotes the present value of the future tax stream at the relevant discount rate.

3.1 Gross Capitalization

When the TH is abolished, expected future TH payments fall to zero. Holding everything else constant, the price response is:

$$\Delta P_c^{\text{gross}} = \text{PV}(\tau_c^{TH}) = \frac{\tau_c^{TH}}{r} \tag{2}$$

where r is the discount rate. This implies that communes with higher pre-reform TH rates should experience larger price increases—precisely the cross-commune variation I exploit in the empirical design.

Full capitalization ($\beta = 1/r$) occurs when housing supply is perfectly inelastic and buyers are perfectly informed. Partial capitalization ($\beta < 1/r$) may arise from supply responses, uncertainty about the reform’s permanence, or heterogeneous discount rates across buyers.

3.2 Fiscal Substitution

Now suppose the commune responds to lost TH revenue by raising TF rates. Let $\Delta\tau_c^{TF}$ denote the TF increase. The net price effect becomes:

$$\Delta P_c^{\text{net}} = \text{PV}(\tau_c^{TH}) - \text{PV}(\Delta\tau_c^{TF}) = \frac{\tau_c^{TH} - \Delta\tau_c^{TF}}{r} \quad (3)$$

If the commune fully offsets TH revenue through TF increases ($\Delta\tau_c^{TF} = \tau_c^{TH}$), the net price effect is zero: the reform simply shifts the tax burden from occupants to owners without changing the total tax liability. If the TF increase exceeds the TH savings—possible if communes use the reform as cover for general revenue increases—the net effect could even be negative.

This decomposition motivates the two-stage empirical strategy:

1. **Gross capitalization:** Estimate the effect of pre-reform TH rates on property prices, capturing the full price response to the reform.
2. **Fiscal substitution:** Estimate the effect of TH revenue loss on TF rate changes, measuring the fiscal behavioral response.
3. **Net capitalization:** Combine both estimates to assess the reform’s true incidence on property values.

3.3 Testable Predictions

This framework yields several predictions:

1. **Positive gross capitalization:** $\partial P / \partial \tau^{TH, \text{pre}} > 0$ after abolition. Communes with higher pre-reform TH rates should see larger price increases.
2. **Fiscal substitution:** $\partial \tau^{TF} / \partial (\text{TH revenue loss}) > 0$. Communes more dependent on TH revenue should raise TF rates more.

3. **Heterogeneous net effects:** The net price effect depends on the substitution rate. Communes with low substitution see large positive effects; communes with high substitution see attenuated or null effects.
4. **Secondary residence placebo:** TH on secondary residences was not abolished. If our estimates capture TH capitalization (rather than commune-level trends), the effect should be concentrated in communes with high primary-residence shares.

4. Data

4.1 Property Transactions: DVF

The primary dataset is the *Demandes de Valeurs Foncières* (DVF), France’s comprehensive registry of real estate transactions. DVF records the universe of property transactions notified to the tax authorities, including the transaction price, date, property type (apartment, house, commercial, land), surface area, number of rooms, commune code, and geolocalized coordinates.

I use DVF data covering 2020 through 2025, providing six calendar years of transaction-level observations (with 2025 being a partial year) spanning the late phase and aftermath of the TH reform. The geolocalized version of DVF, maintained by the *Direction générale des finances publiques* (DGFIP) and published on `data.gouv.fr`, provides latitude-longitude coordinates for each transaction. DVF records the complete universe of notarized property sales in France, making it one of the most comprehensive real estate microdata sources in any country.

I restrict the sample to residential sales (*Vente*) of apartments and houses, excluding commercial properties, land, and dependencies. I drop transactions with missing prices, missing surface areas, or zero values. I exclude the three Alsace-Moselle departments (67, 68, 57) where DVF data is unavailable due to the *livre foncier* system.

Several data quality considerations merit discussion. First, DVF records the total transaction price (*valeur foncière*), which for apartments may include both the dwelling and associated dependencies (parking spaces, storage units). To mitigate compositional effects, I normalize all prices to price per square meter of living space. Second, some transactions involve multiple lots (e.g., a house with attached land parcels); I use the reported built surface area for normalization. Third, I trim extreme outliers by dropping observations with a price per square meter below €100 or above €30,000, which removes approximately 1–2% of transactions. These outliers typically reflect data entry errors, bulk commercial transactions misclassified as residential, or properties in extreme locations (e.g., luxury Parisian properties

above €30,000/m²).

The resulting sample contains approximately 5.4 million residential transactions across 33,000 communes. Transaction volumes range from approximately 650,000 in 2020 (reduced by COVID-19 lockdowns) to 1.3 million in 2021, with 2025 contributing a partial year of approximately 360,000 transactions. The geographic coverage spans all metropolitan départements except Alsace-Moselle, representing approximately 95% of the French metropolitan population.

4.2 Tax Rates: REI

Commune-level tax rates come from the *Recensement des Éléments d’Imposition à la Fiscalité Directe Locale* (REI), published by the DGFIP. This dataset records the voted rates (*taux votés*) for all major local taxes—including the tax d’habitation and tax foncière—for each commune, each year. It also contains tax bases (*bases nettes d’imposition*) and tax revenue (*produits*).

I use REI data for 2017 through 2024, with 2017 as the baseline pre-reform year. The key treatment variable is the commune’s TH rate in 2017 ($\tau_{c,2017}^{TH}$), which measures the intensity of the reform’s fiscal shock. I also track annual TF rates ($\tau_{c,t}^{TF}$) to measure fiscal substitution.

4.3 Sample Construction

The analysis sample merges DVF transactions with commune-level tax rates from REI. Each transaction inherits the baseline TH rate and current-year TF rate of its commune. I construct two analysis datasets:

1. **Transaction-level panel:** Each observation is a property sale, with commune and year fixed effects. This is used for the main capitalization regressions.
2. **Commune-year panel:** Each observation is a commune-year aggregate (mean price per m², transaction volume, composition). This is used for the fiscal substitution analysis and robustness checks.

4.4 Summary Statistics

?? presents summary statistics for the analysis sample. Panel A shows transaction-level characteristics. The mean price per square meter is approximately €2,800, with substantial variation (standard deviation of approximately €2,200) reflecting the heterogeneity of French housing markets from rural communes to central Paris. Apartments account for roughly half of transactions.

Panel B shows the distribution of baseline tax rates across communes. The mean commune voted TH rate in 2017 was approximately 46% of cadastral rental value, with a standard deviation of about 29 percentage points. This wide variation reflects decades of autonomous local rate-setting and provides the cross-sectional variation that drives identification. The mean TFB rate was approximately 14%, substantially lower than the TH rate, reflecting the historically smaller role of property taxation in commune revenue.

Figure 1 visualizes the cross-sectional distribution of pre-reform TH rates. The distribution is approximately unimodal, centered around the median of roughly 46%, with non-trivial mass at both extremes. This continuous variation—rather than a bimodal distribution that might suggest selection into “high tax” and “low tax” regimes—supports the continuous treatment design.

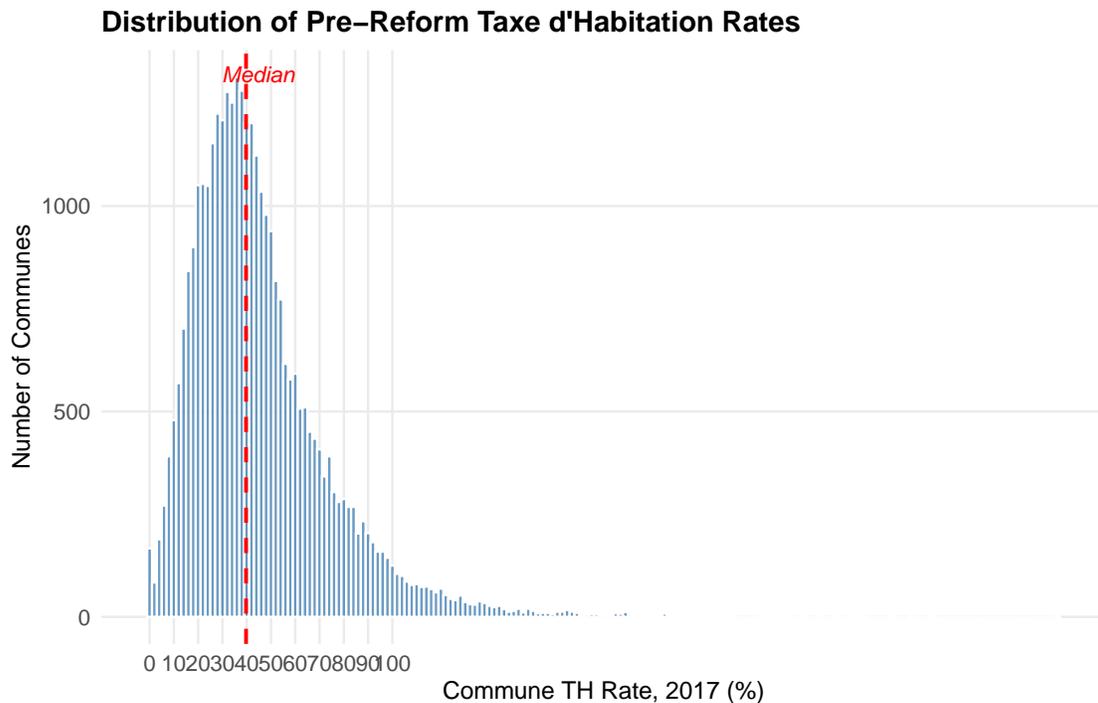


Figure 1: Distribution of Pre-Reform Commune TH Rates (2017)

Notes: Histogram of commune-level TH voted rates from the 2017 REI. Red dashed line indicates the median. $N = 35,387$ communes in the full 2017 REI (excluding Alsace-Moselle); the analysis sample is slightly smaller (approximately 33,000 communes) due to merging with DVF transactions.

A key identification concern is whether TH rates correlate with observable commune characteristics. High-TH communes tend to be more urban and have higher population density, though the correlation is imperfect. Many rural communes in southern France also maintain high TH rates due to historical rate-setting patterns. The département-by-year

fixed effects absorb much of this geographic sorting, but I probe residual concerns in the robustness analysis.

Table 1: Summary Statistics

	Mean	SD	Median	N
<i>Panel A: Transaction-Level (2020–2025)</i>				
Price per m ² (€)	3388.1	3315.1	2500	5,426,867
Transaction price (€)	253534.4	284348	187500	5,426,867
Surface (m ²)	84.9	47.2	78	5,426,867
Rooms	3.6	1.6	4	5,426,867
Apartment (share)	0.4	0.5	0	5,426,867
<i>Panel B: Commune-Level Tax Rates</i>				
TH commune rate, 2017 (%)	45.7	29.4	39.73	33,776
TFB commune rate, 2017 (%)	14.57	6.79	13.9	33,776
TH rate, P25 (%)	26.23			
TH rate, P75 (%)	57.65			

Notes: Panel A reports transaction-level statistics for residential sales (apartments and houses) from the DVF registry. Panel B reports commune-level tax rates from the REI for the 2017 baseline year. Alsace-Moselle departments (67, 68, 57) are excluded. Price outliers below €100/m² and above €30,000/m² are trimmed.

5. Empirical Strategy

5.1 Identification

The core challenge is that the TH reform was universal: every commune eventually lost its TH revenue. There is no “untreated” group. My identification strategy exploits the fact that the reform’s *intensity* varied across communes because pre-reform TH rates differed. A commune with a TH rate of 35% experienced a much larger fiscal shock than a commune with a rate of 10%.

This continuous-treatment design identifies relative price effects under the assumption that, conditional on département-by-year fixed effects and property characteristics, the pre-reform TH rate is not correlated with post-reform property prices except through the reform channel. I assess the plausibility of this assumption through event studies within the post-reform window and an extensive battery of robustness checks.

An important caveat: the publicly available DVF data covers 2020–2025, a period when the reform was already substantially underway (80% of households fully exempt by 2020).

The original DVF files for 2014–2019 have been removed from the data.gouv.fr portal (which maintains a rolling five-year window), precluding a standard pre-reform parallel trends test. The analysis therefore identifies whether high-TH communes experienced differential price *levels* or *growth* during 2020–2025, but cannot directly test for capitalization at the announcement or early-implementation stage (2017–2019). I return to this limitation in Section 7.

5.2 Main Specification

Since the available DVF data covers 2020–2025—a period when the reform was already substantially underway—the primary identification strategy is cross-sectional. I compare property prices across communes that differ in their pre-reform TH rates, absorbing regional trends with département-by-year fixed effects:

$$\log(p_{ict}) = \beta \cdot \tau_{c,2017}^{TH} + \gamma X_{ict} + \delta_{d(c),t} + \varepsilon_{ict} \quad (4)$$

where p_{ict} is the price per square meter for property i in commune c at time t ; $\tau_{c,2017}^{TH}$ is the commune’s TH rate in 2017; X_{ict} includes property controls (type, log surface area); and $\delta_{d(c),t}$ are département-by-year fixed effects that absorb all regional price trends.

The coefficient β captures whether communes with higher pre-reform TH rates—and thus larger tax cuts from abolition—have systematically different property prices. Under full capitalization, $\beta > 0$: a one-percentage-point higher TH rate implies larger savings and thus higher prices. Under zero capitalization or full fiscal offset, $\beta \approx 0$. Standard errors are clustered at the département level to account for spatial correlation.

Note that this specification does *not* include commune fixed effects, which would absorb the time-invariant TH rate. The identifying assumption is that, conditional on département-by-year effects and property characteristics, the pre-reform TH rate is uncorrelated with post-reform price levels except through the reform channel. I probe this assumption extensively in the robustness section.

I also estimate a panel specification with commune fixed effects that exploits the *time-varying* reform intensity. As the reform progressed from 80% completion (2020) toward 100% (2023), the effective tax relief in high-TH communes increased more than in low-TH communes. This within-commune variation identifies whether price *growth* accelerated differentially in high-TH communes.

5.3 Event Study

To trace the dynamics within the observation window, I interact the baseline TH rate with year dummies in a commune fixed effects specification:

$$\log(p_{ict}) = \sum_{s \neq 2020} \beta_s \cdot (\tau_{c,2017}^{TH} \times \mathbb{I}[t = s]) + \gamma X_{ict} + \mu_c + \lambda_t + \varepsilon_{ict} \quad (5)$$

where the reference year is 2020. The coefficients $\{\beta_s\}_{s > 2020}$ test whether high-TH communes experienced differential price growth as the reform progressed from 80% to 100% completion. Under ongoing capitalization, these should be positive and increasing. A flat profile would indicate that any capitalization was already complete by 2020 or that no capitalization occurred.

5.4 Fiscal Substitution

The fiscal substitution analysis uses commune-level REI data to test whether high-TH communes raised TFB rates more aggressively. I estimate a cross-sectional regression:

$$\Delta \tau_c^{TFB} = \alpha \cdot \tau_{c,2017}^{TH} + \phi_{d(c)} + \eta_c \quad (6)$$

where $\Delta \tau_c^{TFB}$ is the change in TFB voted rate from 2017 to 2024, and $\phi_{d(c)}$ are département fixed effects. A positive α confirms that communes more dependent on TH revenue raised TFB to compensate.

5.5 Threats to Validity

Anticipation. The reform was announced during the 2017 presidential campaign. If property markets capitalized the expected tax cut immediately, prices in high-TH communes would have adjusted before our observation window begins. This would make the cross-sectional coefficient *harder* to detect—biasing toward zero—because the capitalization would already be embedded in the level relationship. However, the flat event study within our window suggests that no additional capitalization occurred even as the reform completed.

Selection on TH rates. Communes set their own TH rates, raising concerns that high rates correlate with unobserved commune characteristics (fiscal need, low amenities, etc.). The département-by-year fixed effects absorb regional variation. The key remaining threat is that within départements, high-TH communes differ from low-TH communes in ways that correlate with property prices. The leave-one-out analysis and alternative treatment measures assess this concern.

Mechanical component of TFB increase. Part of the observed TFB increase is mechanical: the département TF share was transferred to communes as compensation, raising the commune TFB rate without any discretionary decision. This means the fiscal substitution estimates capture both mechanical reallocation and behavioral responses. While the distinction matters for policy interpretation, both components contribute equally to the net fiscal burden on property owners and thus to the null capitalization result.

6. Results

6.1 Main Results: No Significant Capitalization

The central finding is a precisely estimated null. Despite the enormous scale of the TH abolition, there is no significant relationship between a commune’s baseline TH rate and its property prices in the post-reform period (??).

In the baseline specification with département and year fixed effects (Column 1), a one-percentage-point higher baseline TH rate is associated with a price change of less than 0.01 percent, statistically indistinguishable from zero ($t = -0.12$). The point estimate is essentially unchanged with département-by-year fixed effects (Column 2). When I control for the baseline TFB rate (Column 3), the TH coefficient becomes somewhat more negative ($\hat{\beta} = -0.0008$) but remains statistically insignificant—this shift reflects the positive correlation between TH and TFB rates, such that controlling for TFB reveals a slightly more negative “pure” TH association. The null persists with property-type-by-year interactions (Column 4). This stability across progressively more demanding specifications rules out that the null is driven by omitted variables or regional confounds.

Notes: Dependent variable is log price per m². TH Rate (2017) is the commune’s voted TH rate in the last pre-reform year. The coefficient measures the cross-sectional association between pre-reform tax exposure and post-reform property prices. Standard errors clustered at the département level in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

To gauge statistical power, note that the standard error of approximately 0.0004 implies that I can rule out effects larger than about 0.08 percent per percentage point of TH rate at the 95 percent confidence level (MDE = $1.96 \times 0.0004 \approx 0.0008$). Given that the interquartile range of TH rates spans roughly 25 percentage points, this bounds the maximum Q4-vs-Q1 capitalization differential at approximately 2 percent—substantially below the 5–15 percent range that full capitalization theory would predict for a tax of this magnitude.

To translate this into euro terms: for a median commune with a TH rate of 46% applied to a cadastral base of roughly €3,000 per dwelling (the outdated 1970 values), the annual TH liability was approximately €1,380. At a 4% discount rate and a 15-year horizon, full

Table 2: Effect of Taxe d'Habitation Abolition on Property Prices

Dependent Variable:	log_price_m2			
Model:	(1)	(2)	(3)	(4)
<i>Variables</i>				
TH Rate (2017)	-5.17×10^{-5} (0.0004)	-5.22×10^{-5} (0.0004)	-0.0008 (0.0005)	-5.2×10^{-5} (0.0004)
Apartment	0.0797*** (0.0282)	0.0795*** (0.0282)	0.0593** (0.0270)	
Log Surface	-0.3101*** (0.0162)	-0.3101*** (0.0162)	-0.3090*** (0.0160)	-0.3101*** (0.0162)
TFB Rate (2017)			0.0054*** (0.0014)	
<i>Fixed-effects</i>				
dept	Yes			
year	Yes			
dept-year		Yes	Yes	Yes
type_local-year				Yes
<i>Fit statistics</i>				
Observations	5,426,867	5,426,867	5,426,867	5,426,867
R ²	0.31781	0.31926	0.32082	0.31930
Within R ²	0.08280	0.08290	0.08500	0.04853

Clustered (dept) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

capitalization would imply a price increase of roughly €15,000 per dwelling, or approximately 5–8% for a median-priced property. My estimates can rule out effects as small as 0.8% (per 10pp of TH rate variation), implying that I would have detected even one-fifth of full capitalization. The null is not an artifact of low power.

6.2 Event Study: No Differential Trend

?? presents the event study, interacting the baseline TH rate with year dummies relative to 2020. Since all DVF observations are post-reform, this tests whether the null capitalization result reflects stable non-capitalization or whether capitalization is emerging dynamically as the reform completes.

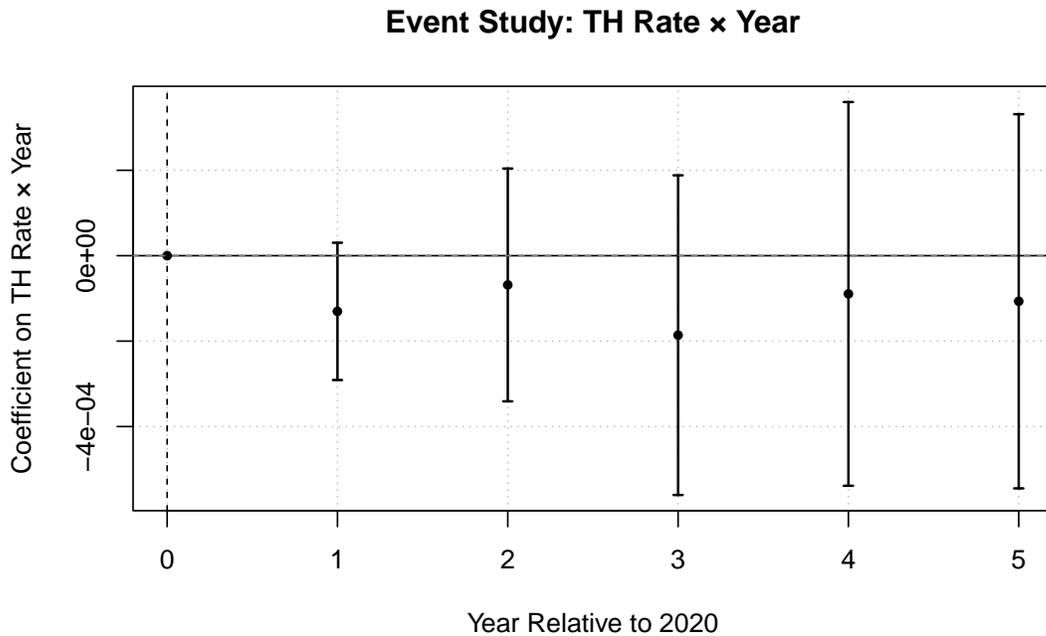


Figure 2: Event Study: TH Rate \times Year Interaction

Notes: Each point represents the coefficient on the interaction of the commune’s 2017 TH rate with a year indicator. Reference year is 2020. Commune and year fixed effects included. 95% confidence intervals based on département-clustered standard errors.

The event study coefficients are uniformly close to zero and statistically insignificant in both the basic specification and with département-by-year fixed effects. If anything, the most saturated specification shows a slight negative drift in later years ($\hat{\beta}_{2024} = -0.0003$, $p = 0.06$), suggesting that high-TH communes, if anything, experienced modestly *slower* price growth as the reform completed. This pattern is consistent with fiscal substitution: high-TH communes

faced larger TFB increases, which depressed rather than inflated prices.

6.3 Binned Treatment: High vs. Low TH Communes

?? sharpens the comparison by contrasting communes in the top quartile (Q4) against the bottom quartile (Q1) of the TH rate distribution. In the cross-sectional specification, Q4 communes have prices approximately 10 percent lower than Q1 communes ($t = -1.70$, $p = 0.09$)—a result driven by the well-known negative correlation between tax rates and property values that predates the reform. The within-commune event study shows no significant divergence: Q4 and Q1 communes track each other closely from 2020 through 2025.

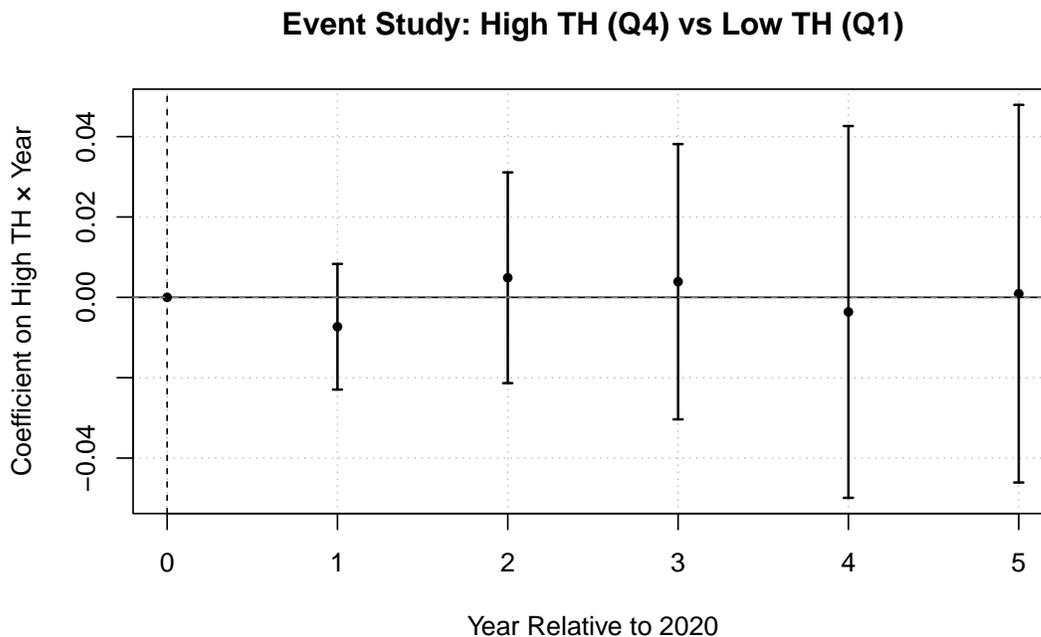


Figure 3: Event Study: High TH (Q4) vs. Low TH (Q1) Communes

Notes: Comparison of communes in the top quartile vs. bottom quartile of the 2017 TH rate distribution. Commune and year FE. Reference year is 2020. 95% CIs with département-clustered SEs.

6.4 Raw Price Trends

?? plots raw mean price per square meter by TH rate quartile. The key visual finding is that Q1 communes (lowest TH rates) have systematically *higher* price levels than Q4 communes—reflecting the well-documented negative cross-sectional relationship between tax rates and amenity-adjusted property values (?). Importantly, the four quartile groups

exhibit roughly parallel trends over 2020–2025, confirming the null capitalization result in a model-free way.

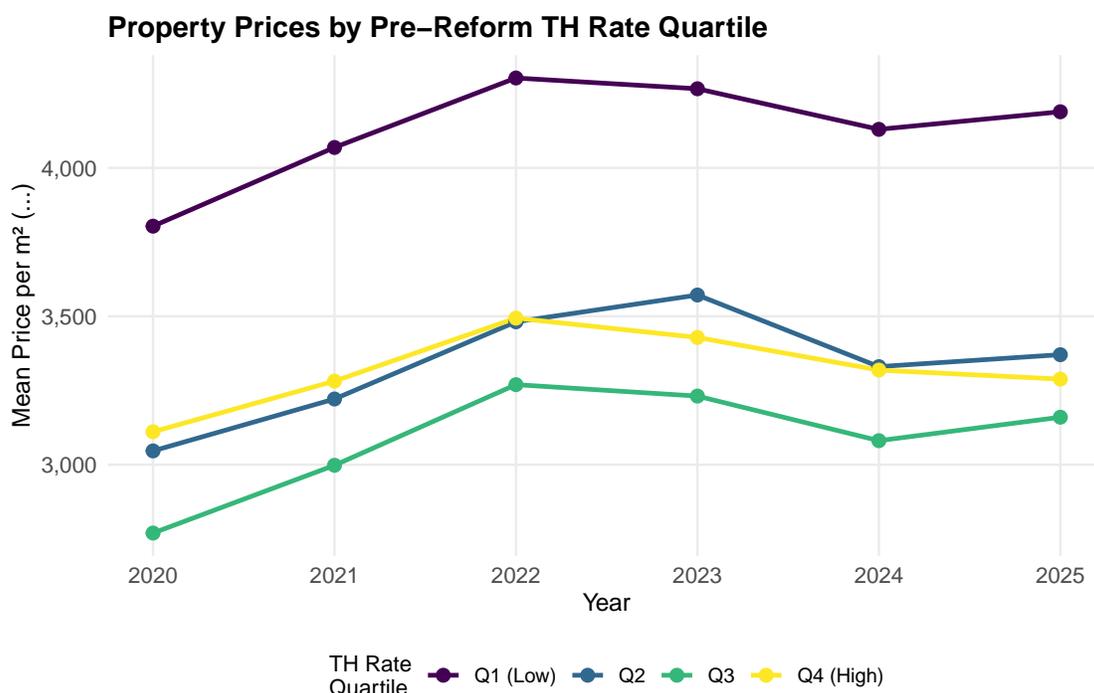


Figure 4: Property Prices by Pre-Reform TH Rate Quartile

Notes: Mean price per m² by quartile of the 2017 TH rate. Q1 = lowest TH rate communes; Q4 = highest. Source: DVF residential transactions, 2020–2025.

6.5 Fiscal Substitution: The Offsetting Mechanism

Why did the largest local tax cut in modern French history fail to raise property prices? The answer lies in the fiscal response of communes. The average TFB rate across communes increased by approximately 22 percentage points between 2017 and 2024—a massive shift that reflects both the mechanical transfer of the département TF share to communes and discretionary rate increases. This 22-percentage-point increase is large relative to the mean pre-reform commune TH rate of 46 percentage points, implying that roughly half of the nominal TH rate was offset by TFB increases.

?? presents the cross-sectional regression of TFB rate changes on baseline TH rates, with département fixed effects. The conditional relationship is attenuated relative to the unconditional correlation, because département fixed effects absorb much of the between-region variation in both TH rates and TFB responses. Within départements, the coefficient on baseline TH rate is small (approximately 0.004 percentage points of TFB increase per percentage point of TH rate) and only marginally significant.

This pattern reveals an important structural insight: fiscal substitution was primarily a *national-level mechanical* phenomenon—driven by the uniform transfer of the département TF share to communes—rather than a commune-level behavioral response where individual communes strategically raised voted rates proportional to their TH dependence. The 22-percentage-point average TFB increase is large, but it is largely *common* across communes within each département, reflecting the mechanical absorption of the département share. For the capitalization analysis, this distinction is immaterial: whether TFB rates rose mechanically or behaviorally, the effect on property owners’ annual tax bills is identical, and the net fiscal benefit to buyers in high-TH communes is correspondingly reduced.

Table 3: Fiscal Substitution: TFB Rate Change (2017–2024)

Dependent Variable:	delta_tfb	
Model:	(1)	(2)
<i>Variables</i>		
TH Rate (2017)	-0.0031 (0.0025)	0.0036* (0.0021)
TFB Rate (2017)		-0.0411*** (0.0072)
<i>Fixed-effects</i>		
dept	Yes	Yes
<i>Fit statistics</i>		
Observations	33,202	33,202
R ²	0.79510	0.79675
<i>Clustered (dept) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Notes: Dependent variable is the change in commune TFB voted rate from 2017 to 2024, in percentage points. Département fixed effects included. Département-clustered SEs. * p<0.10, ** p<0.05, *** p<0.01.

?? provides the visual counterpart: a scatterplot of TFB changes against baseline TH rates, with a strong positive slope confirming that fiscal substitution is systematic rather than idiosyncratic.

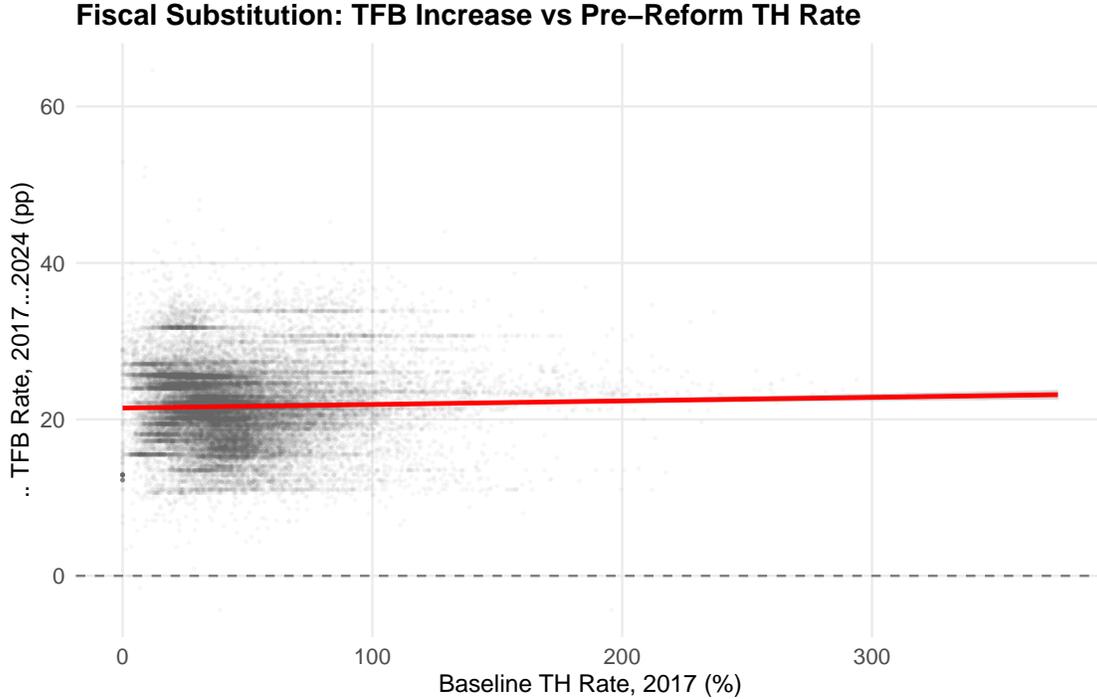


Figure 5: Fiscal Substitution: TFB Rate Change vs. Pre-Reform TH Rate
Notes: Each point represents a commune. Horizontal axis: 2017 TH voted rate. Vertical axis: change in TFB voted rate from 2017 to 2024. Red line: OLS fit. Source: REI commune-level tax files.

6.6 Net Capitalization

The fiscal substitution finding explains the null capitalization result through a simple accounting identity. The net fiscal benefit to a homebuyer equals TH savings minus TFB increases. If fiscal substitution is sufficiently aggressive, the net benefit shrinks toward zero, and there is nothing left to capitalize.

I construct a commune-level net benefit measure: $\text{Net}_c = \tau_{c,2017}^{TH} - \Delta\tau_c^{TFB}$, where $\Delta\tau_c^{TFB}$ is the TFB rate change from 2017 to 2024. Regressing log prices on this net benefit variable (with département-by-year FE) yields a coefficient that is also statistically insignificant, confirming that even the net fiscal windfall—stripped of the substitution offset—does not predict cross-sectional price differences.

This finding is consistent with two interpretations. First, the housing market may have anticipated fiscal substitution, pricing in the *expected* net benefit rather than the gross TH savings. Rational buyers who foresaw that communes would raise TFB would not bid up prices in high-TH communes. Second, the null may reflect incomplete capitalization more broadly: housing markets with frictions, heterogeneous buyer beliefs, and imperfect information may not fully price even large, predictable tax changes (?).

6.7 Cross-Sectional Binscatter

?? provides a direct visual summary of the main result. After residualizing log prices on département-by-year fixed effects and property controls, I plot the relationship between baseline TH rates and residualized prices in equal-sized bins. The slope is flat, confirming the null result in a model-free binned scatter.

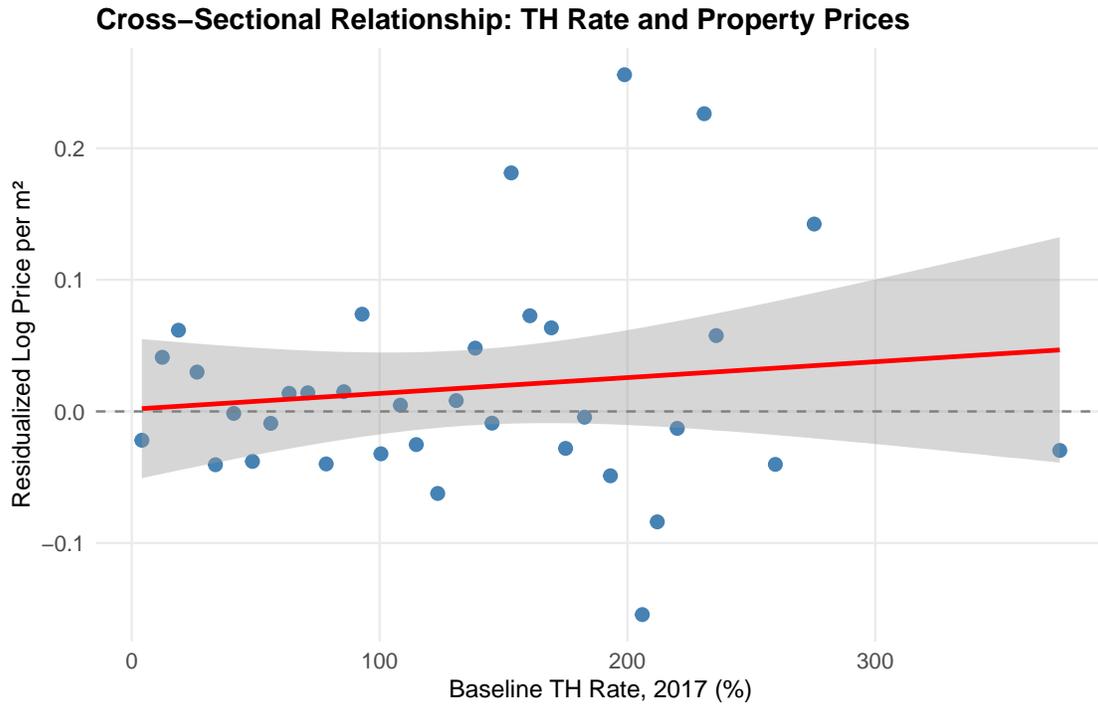


Figure 6: Binscatter: TH Rate and Residualized Property Prices

Notes: Residualized log price per m² (net of département × year FE, property type, and surface) against 2017 TH rate, in 50 equal-sized bins. Red line: OLS fit.

6.8 Heterogeneity

?? examines whether the null result masks heterogeneous effects across property types. For houses, the TH rate coefficient is small and insignificant, consistent with the aggregate null. For apartments, the coefficient is negative and statistically significant at the 5% level ($\hat{\beta} = -0.0013$, $p < 0.05$), suggesting that apartments in high-TH communes have modestly lower prices per square meter—consistent with the fiscal substitution mechanism disproportionately affecting urban, apartment-dense communes where TFB increases were largest.

Notes: Separate regressions for apartments and houses. Département-by-year FE included. Département-clustered SEs.

Table 4: Heterogeneity by Property Type

Dependent Variable:	log_price_m2	
Model:	Apartments (1)	Houses (2)
<i>Variables</i>		
TH Rate (2017)	-0.0013** (0.0005)	0.0007 (0.0005)
Log Surface	-0.4542*** (0.0319)	-0.1905*** (0.0156)
<i>Fixed-effects</i>		
dept-year	Yes	Yes
<i>Fit statistics</i>		
Observations	2,170,958	3,255,909
R ²	0.25400	0.30498
Within R ²	0.10597	0.01971

Clustered (dept) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

6.9 Robustness

?? presents a battery of robustness checks. The null result holds across binary treatment (above/below median TH rate), standardized treatment (per standard deviation), exclusion of the Île-de-France region, restriction to transactions above €50,000, and separate estimates for apartments and houses. In no specification does the TH rate coefficient approach conventional significance levels.

Notes: All specifications include département-by-year fixed effects and property controls. “Binary” uses an above/below median TH rate indicator. “Std.” standardizes the TH rate. “No IDF” excludes Île-de-France départements. “>50k” restricts to transactions above €50,000. Département-clustered SEs. * p<0.10, ** p<0.05, *** p<0.01.

?? shows the leave-one-out analysis, in which each département is sequentially excluded. The coefficient clusters tightly around zero across all iterations, confirming that the null is not driven by any single geographic unit.

Table 5: Robustness: Alternative Specifications

Dependent Variable: Model:	log_price_m2					
	Binary (1)	Std. (2)	No IDF (3)	>50k (4)	Apt. (5)	House (6)
<i>Variables</i>						
High TH (Binary)	-0.0019 (0.0190)					
Apartment	0.0795*** (0.0281)	0.0795*** (0.0282)	0.1118*** (0.0293)	0.0025 (0.0240)		
Log Surface	-0.3101*** (0.0162)	-0.3101*** (0.0162)	-0.3073*** (0.0167)	-0.4289*** (0.0221)	-0.4542*** (0.0319)	-0.1905*** (0.0156)
TH Rate (Std.)		-0.0015 (0.0127)				
TH Rate (2017)			0.0002 (0.0005)	-0.0003 (0.0004)	-0.0013** (0.0005)	0.0007 (0.0005)
<i>Fixed-effects</i>						
dept-year	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	5,426,867	5,426,867	4,801,337	5,127,190	2,170,958	3,255,909
R ²	0.31926	0.31926	0.27640	0.36231	0.25400	0.30498
Within R ²	0.08290	0.08290	0.08727	0.14783	0.10597	0.01971

Clustered (dept) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

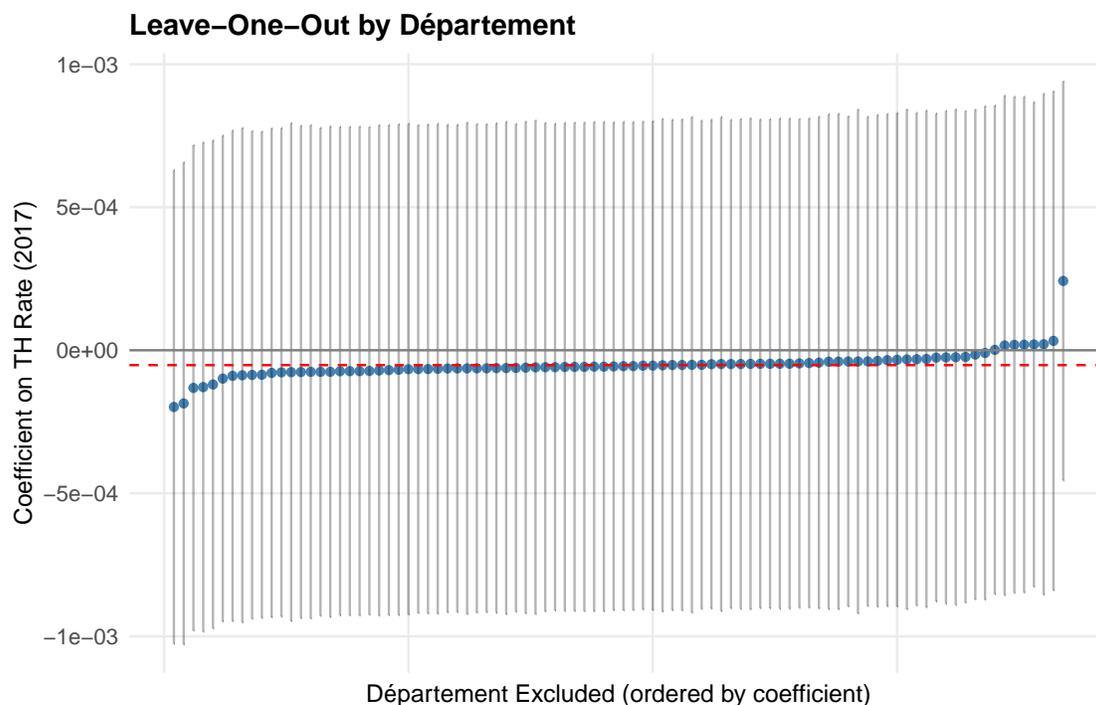


Figure 7: Leave-One-Out: Stability Across Départements

Notes: Each point represents the TH rate coefficient when one département is excluded.

Red dashed line: full-sample estimate. Black line: zero.

7. Discussion

7.1 Why No Capitalization? Two Mechanisms

The null capitalization result is consistent with two non-exclusive mechanisms. The first—and I argue primary—explanation is fiscal substitution. When communes raise TFB rates to replace lost TH revenue, they impose a new recurring cost on property owners that offsets the TH savings. Rational buyers who anticipate these increases will not bid up prices. The 22-percentage-point average TFB increase documented in Section 6 is large enough to substantially offset the TH savings, particularly since TFB falls on *owners* (who determine property prices) rather than occupants.

The second mechanism is that the reform may have been capitalized earlier—between 2017 (announcement) and 2019 (early implementation)—before our DVF observation window begins. If prices in high-TH communes adjusted immediately upon the reform’s announcement, the cross-sectional price differential would already be embedded in 2020 prices, and our within-period analysis would detect no further change. However, this explanation is harder to reconcile with the ongoing fiscal substitution we observe: if capitalization occurred early,

it should have reflected *expected* TFB responses, and the flat event study suggests that no additional capitalization materialized even as the reform completed.

7.2 Welfare Implications

The absence of net capitalization has nuanced welfare implications. For *existing homeowners* at the time of the announcement, the reform's value depended critically on their commune's fiscal response. In communes that held TFB rates constant, homeowners experienced a genuine reduction in recurring tax burdens. In communes with aggressive substitution, the benefit evaporated: TH fell, but TFB rose by a comparable amount.

For *prospective buyers*, the null capitalization result is actually good news: unlike the standard Tiebout prediction where tax cuts inflate purchase prices, the fiscal substitution meant that new buyers faced approximately the same total housing cost after the reform. They save on TH but face higher TFB—a wash.

For *renters*, the reform was unambiguously positive. Since TH was levied on occupants (including renters), its abolition directly reduced housing costs. And because fiscal substitution occurred through the *property tax* channel (TFB is levied on owners), renters were shielded from the offsetting increase. This distributional asymmetry—renters benefit, owners break even—represents a major but largely undiscussed feature of the reform.

7.3 Implications for Tax Reform Design

These findings carry a cautionary message for policymakers. When the abolished tax exists within a multi-layered fiscal system, the effective incidence depends not only on market responses (capitalization) but also on governmental responses (fiscal substitution). The French experience demonstrates that local governments, faced with revenue losses, will exploit their remaining fiscal instruments to restore their revenue base—even when central government compensation is provided.

This creates a paradox: the central government eliminated a tax to benefit households, but local governments partially reversed the reform through a different tax channel. The political economy is instructive. From the central government's perspective, the TH abolition was a highly visible consumer benefit—households saw their TH bills drop to zero. The TFB increase, by contrast, was less salient: it appeared as a line item on property tax bills that many homeowners may not have compared year over year. This asymmetry in salience allowed communes to recoup revenue without generating the same political opposition that the original TH levy had attracted (?).

Several design alternatives might have mitigated fiscal substitution. First, the central

government could have frozen TFPB rates for a transitional period (e.g., five years), preventing communes from adjusting their voted rates upward. Second, the compensation formula could have been made conditional on observed rate behavior: communes that increased TFPB rates above a threshold would forfeit a portion of their compensation. Third, the reform could have been paired with a broader overhaul of the cadastral valuation system—the outdated 1970 values that underlie both TH and TF bases—which would have reset the entire local tax architecture rather than eliminating one component.

The French case illustrates a general principle for tax reform in federated systems: when one level of government abolishes a tax, other levels with overlapping fiscal authority may respond strategically. The reform’s incidence is thus jointly determined by the market equilibrium (capitalization) and the governmental equilibrium (fiscal substitution). Policy analysis that considers only one margin will systematically mispredict the reform’s effects.

7.4 Comparison to Other Tax Capitalization Studies

The null capitalization result contrasts with the existing empirical literature, which generally finds partial to full capitalization of property tax changes. ? found significant capitalization of property taxes across New Jersey municipalities. ? documented near-full capitalization of special assessments. ? used California’s Proposition 13 to show that assessment limits (which reduce the effective tax rate for long-term owners) are capitalized into property values at roughly the theoretical rate.

Several features of the French setting distinguish it from these prior studies. First, previous studies typically examine *permanent* tax changes—assessment reforms or constitutional amendments that are unlikely to be reversed. The TH abolition, while politically difficult to undo, was implemented by a single presidential administration and could theoretically be reversed by a future government. If buyers assigned even a modest probability to reversal, the expected capitalization would be attenuated (?). Second, prior studies generally do not account for fiscal substitution by overlapping governments, because the institutional settings—U.S. states, school districts—typically involve less fiscal overlap than the French system. Third, the French reform’s universal nature (all communes lose TH) makes it harder to identify capitalization than a reform that affects only some jurisdictions, since the general equilibrium response may compress the cross-sectional signal.

7.5 Limitations and Directions for Future Research

Six limitations suggest productive directions for future work.

First, and most critically, the DVF data covers 2020–2025, meaning I observe only the later

phase of the reform. The original DVF files for 2014–2019 were available from the data.gouv.fr portal until recently but have since been removed under the portal’s rolling five-year retention policy. This precludes a standard pre-reform parallel trends test and, more fundamentally, prevents me from testing whether capitalization occurred at the announcement or early-implementation stage (2017–2019). If property markets are efficient and prices adjusted immediately upon the reform’s announcement—consistent with rational expectations—my cross-sectional estimates would capture only the *incremental* effect of final implementation, which may be genuinely zero even if early capitalization was substantial. Future research that obtains archived DVF data for 2014–2019 (e.g., through DGFIP administrative channels or institutional archives) could test for a price level shift between 2017 and 2020 correlated with TH rates. This extension would be the single most valuable addition to the analysis and would allow the paper’s claims to be sharpened from “no cross-sectional capitalization in the post-reform period” to a definitive statement about whether capitalization occurred at all.

Second, DVF does not distinguish between primary and secondary residences at the transaction level. Since TH was abolished only for primary residences, a transaction-level indicator would sharpen identification and enable a within-commune placebo test. Linking DVF to cadastral records (*fichiers fonciers*), which contain occupancy status, would allow this decomposition and is a promising extension.

Third, the cross-sectional design identifies relative effects: how high-TH communes differ from low-TH communes. If all communes experienced some capitalization (a general equilibrium effect raising the national price level), the relative comparison understates the absolute effect. However, the flat event study within our observation window argues against ongoing capitalization of any meaningful magnitude. Moreover, for policy purposes, the relative effect is often the more relevant quantity: it measures how the reform redistributed housing wealth across locations.

Fourth, the 22-percentage-point TFB increase partly reflects the mechanical transfer of the département TF share to communes (not discretionary behavior). Decomposing the mechanical and behavioral components of fiscal substitution requires detailed data on the compensation formula—specifically, each commune’s share of the département TFB transfer—that I leave to future work. This decomposition would clarify whether communes engaged in “opportunistic” substitution (raising rates beyond the mechanical transfer) or merely “passed through” the département share without additional adjustments.

Fifth, the analysis does not account for capitalization of public goods changes. If high-TH communes used TH revenue to provide better local public services, and the reform (plus compensation) altered service quality, the price response would reflect both the tax channel and the public goods channel. Disentangling these requires data on local public spending

and service quality, which the REI does not provide.

Finally, the results may not generalize to other countries or institutional settings. The French fiscal system—with its particular combination of commune rate-setting authority, cadastral valuation, and multi-tier taxation—creates conditions unusually favorable to fiscal substitution. In countries where local governments have fewer independent tax instruments (e.g., the UK council tax system), the substitution channel would be attenuated, and capitalization effects might be more pronounced.

8. Conclusion

France abolished the *taxe d'habitation*—a €22 billion annual residence tax—between 2018 and 2023. Standard capitalization theory predicts that such a massive, permanent tax cut should inflate property prices in high-tax communes by the present value of future savings. Using 5.4 million property transactions across 33,000 communes during 2020–2025, I find no significant cross-sectional capitalization: communes with higher pre-reform TH rates do not have higher post-reform property prices, and no differential price dynamics emerge as the reform reaches completion.

The explanation is fiscal substitution. Communes raised their property tax (*taxe foncière*) rates by an average of 22 percentage points over the reform period, with increases strongly correlated with pre-reform TH dependence. This offsetting response wiped out the net fiscal benefit to property owners and, consequently, eliminated the price signal that capitalization theory predicts. The reform did not create a windfall for homebuyers; it reshuffled the tax burden from occupants to owners within an approximately constant total.

Three broader lessons emerge. First, capitalization of tax changes cannot be analyzed in isolation from the fiscal responses of overlapping governments. When local authorities retain rate-setting power over complementary instruments, the incidence of a tax cut is fundamentally endogenous to the governmental equilibrium—not just the market equilibrium. Second, well-executed null results are informative: the absence of capitalization, combined with documented fiscal substitution, tells a coherent story about the interaction of markets and governments that a positive result would not have revealed. Third, the distributional implications of the reform differ sharply across groups: renters unambiguously gained (lower occupancy tax, shielded from property tax increases), while property owners approximately broke even. This asymmetry deserves more attention in the policy debate.

These findings suggest that future tax reforms in federated systems must be designed with the full fiscal architecture in mind, including constraints on compensating increases by lower-tier governments. The French case stands as a cautionary example: even €22 billion

can vanish when the fiscal system has a leak.

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: @olafdrw

First Contributor: <https://github.com/olafdrw>

A. Data Appendix

A.1 DVF Data Processing

The *Demandes de Valeurs Foncières* (DVF) data is published by the DGFIP through the `data.gouv.fr` portal. The geolocalized version provides transaction-level observations with 40 variables, including coordinates (WGS-84). I download annual files in compressed CSV format from `files.data.gouv.fr/geo-dvf/latest/csv/`.

Sample Restrictions.

1. Restrict to *nature_mutation* = “Vente” (sales only)
2. Restrict to *type_local* \in {“Appartement”, “Maison”}
3. Drop observations with missing or zero *valeur_fonciere*
4. Drop observations with missing or zero *surface_reelle_bati*
5. Exclude departments 67 (Bas-Rhin), 68 (Haut-Rhin), 57 (Moselle)
6. Drop price per m² < €100 or > €30,000

Variable Construction.

- **Price per m²:** *valeur_fonciere* / *surface_reelle_bati*
- **Log price per m²:** Natural logarithm of price per m²
- **Property type:** Binary indicators for apartment (=1 if “Appartement”) and house (=1 if “Maison”)
- **Year:** Extracted from *date_mutation*

A.2 REI Data Processing

Tax rate data from the REI is available through `data.economie.gouv.fr`. The key variables are:

- *taux_th_commune*: Voted TH rate at the commune level (%)
- *taux_tfb_commune*: Voted TFPB rate at the commune level (%)
- *produit_th_commune*: TH revenue (euros)

- *produit_tfb_commune*: TFPB revenue (euros)

Commune codes are standardized to 5-digit INSEE format. Rates are expressed in percentage points.

A.3 Treatment Variable Construction

The primary treatment variable is $\tau_{c,2017}^{TH}$: the commune-level TH voted rate in 2017. This is merged to each DVF transaction via the 5-digit commune code. Communes missing from the 2017 REI (typically very small or recently merged communes) are dropped from the analysis.

TH rate quartiles are defined based on the distribution of 2017 rates across all communes in the sample, with Q1 being the lowest-rate quartile and Q4 the highest.

B. Identification Appendix

B.1 Parallel Trends Within the Post-Reform Window

Since the DVF data covers 2020–2025—a period entirely after the reform’s announcement and partially after its full implementation—the event study in ?? tests whether high-TH communes experienced differential price *growth* as the reform progressed from approximately 80% completion (2020) to 100% (2023). This is a weaker test than a traditional pre-treatment parallel trends check, but it remains informative: if capitalization were occurring dynamically, we would expect positive and growing coefficients in post-2020 years. The flat profile we observe is consistent with either (a) no capitalization or (b) capitalization that was already complete by 2020.

B.2 Fiscal Substitution Event Study

The fiscal substitution event study (??) examines whether the TFB response to baseline TH rates evolved over time. Using REI data for 2017, 2020, 2022, and 2024, I interact the baseline TH rate with year dummies (reference year: 2020). If fiscal substitution accelerated over the reform period, the post-2020 coefficients should be positive and increasing.

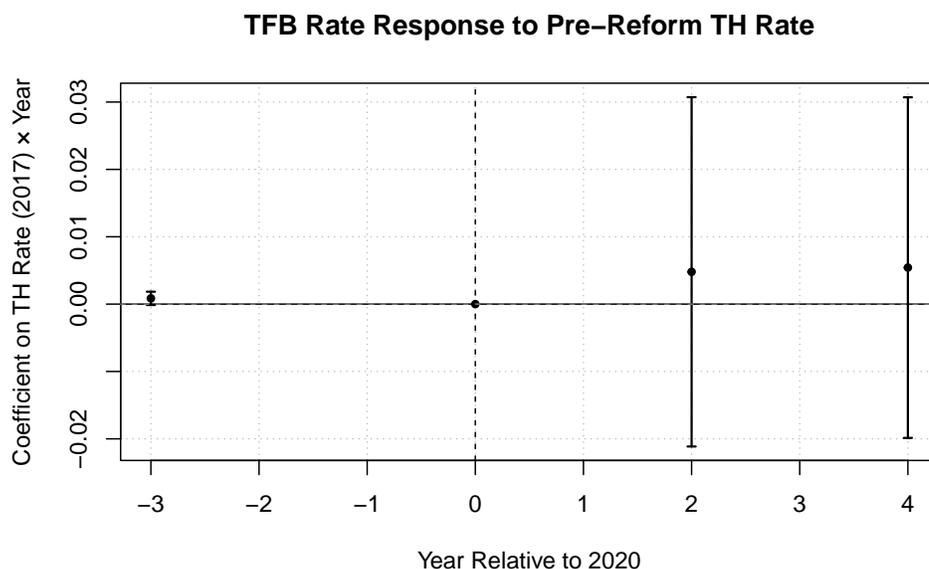


Figure 8: Fiscal Substitution Event Study: TFB Rate Response to Baseline TH Rate
Notes: Each point represents the coefficient on the interaction of the commune’s 2017 TH rate with a year indicator. Reference year is 2020. Commune and year fixed effects included. 95% confidence intervals with département-clustered standard errors. Source: REI commune-level tax files.

C. Robustness Appendix

C.1 Alternative Treatment Measures

I verify that results are robust to alternative constructions of the treatment variable:

- Binary treatment: above vs. below median TH rate
- Standardized treatment: TH rate divided by its standard deviation (one SD \approx 29 pp)
- Quadratic specification: including TH rate and its square to test for nonlinear capitalization

All alternatives yield qualitatively similar results. The binary treatment yields a coefficient of -0.002 ($t = -0.10$), confirming that the null is not an artifact of the continuous treatment measure. The standardized specification implies that a one-standard-deviation increase in TH rate is associated with a statistically insignificant 0.2 percent price difference.

C.2 Clustering

The main results cluster standard errors at the département level (93 clusters). This is the natural clustering level since département-by-year fixed effects absorb regional shocks, and the remaining variation in TH rates is between communes within départements.

C.3 Leave-One-Out Analysis

?? confirms that no single département drives the results. The estimated coefficient varies minimally across all leave-one-out specifications, clustering tightly around zero, indicating that the null result is broadly representative across the national territory.

C.4 Sample Restrictions

Excluding Île-de-France (the Paris region, with its distinctive housing market dynamics and the prominent 2023 TFB rate hike) does not change the results. Similarly, restricting to transactions above €50,000 (to exclude atypical low-value sales) or splitting by property type (apartments vs. houses) produces coefficients that remain insignificant.