Who Captures a Tax Cut? Property Price Capitalization from France's Taxe d'Habitation Abolition

APEP Autonomous Research*

March 3, 2026

Abstract

When a government hands 23.4 billion back to its citizens, does the money stay in their pockets? I study France's phased abolition of the *taxe d'habitation* (TH), exploiting cross-commune variation in pre-reform tax rates as a continuous treatment dose in a difference-in-differences framework. Under baseline controls, a one-standard-deviation increase in TH exposure raises apartment prices per square meter by 2.3 percent ($p < 0.001$), with the effect absent for houses. However, this result is sensitive: adding département × year fixed effects eliminates the apartment effect, and the capitalization is concentrated in years after a data-source transition. The apartment event study shows clean pre-trends ($p = 0.62$) and a pattern consistent with delayed capitalization, but the evidence does not survive the most demanding specifications. These findings suggest that identifying property tax capitalization requires careful attention to spatial trends and data construction, even in settings with clean institutional variation.

JEL Codes: H22, H71, R21, R31

Keywords: tax capitalization, property tax, taxe d'habitation, housing prices, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

When a government abolishes a 23.4 billion annual property tax, who actually benefits? Between 2018 and 2023, France eliminated the *taxe d'habitation* (TH) on primary residences—a tax paid by every household in every commune—with the explicit goal of boosting purchasing power. In a world of fixed housing supply, however, economic theory predicts a more cynical outcome: the tax cut is capitalized into higher property prices, enriching current owners at the expense of future buyers (**?**). Whether this prediction holds for a reform of this scale, affecting over 30 million households, has direct implications for how we evaluate one of France's largest fiscal experiments.

This paper provides the first causal estimate of property price capitalization from the TH abolition. The empirical challenge is straightforward: the reform applied to all French communes simultaneously, eliminating simple before-after comparisons. I exploit the fact that communes varied widely in their pre-reform TH rates—from near zero to over 26 percent—generating cross-sectional variation in the size of the tax relief. Communes with higher TH rates experienced larger reductions in the tax burden on occupants, providing a continuous "dose" of treatment intensity. I combine this cross-commune variation with the timing of the reform in a difference-in-differences framework, comparing price trajectories in high-TH versus low-TH communes before and after 2018.

The data initially reveal a tale of two markets. Under baseline commune and year fixed effects, apartment prices per square meter rise by 2.3 percent for every standard deviation of tax relief ($p < 0.001$), while houses show no response. However, this apparent capitalization does not survive the most stringent controls: with département $\times$ year fixed effects, the apartment effect vanishes (-0.006, $p = 0.38$), and restricting to the pre-2021 data source yields a null result. The apartment-specific event study offers a more nuanced picture: pre-trends are clean (joint F-test $p = 0.62$), and post-reform coefficients grow steadily to 4–5 percent by 2022–2024. But these late-period effects coincide with a change in data source (from aggregate CdD statistics to transaction-level geo-DVF data), complicating causal interpretation.

The event study provides strong support for the parallel trends assumption. Pre-reform coefficients for 2014–2016 are small and statistically insignificant, centered around zero. The anticipation test finds a marginally significant negative coefficient for 2017—the year the reform was announced—but in the wrong direction, suggesting no meaningful anticipation effects on prices. The leave-one-out analysis across 93 départements shows that no single département drives the result, with the coefficient ranging from -0.001 to 0.004.

The heterogeneity results are suggestive about mechanisms but must be interpreted cautiously given the specification sensitivity documented above. Under baseline controls,

the apartment-versus-house contrast (2.3% versus 0.01%, not significant) suggests that capitalization operates most strongly in thicker, more standardized markets (**?**). However, since the apartment effect is not robust to département × year fixed effects, the contrast may partly reflect differential spatial trends across property types. Supply-side heterogeneity provides a complementary test: theory predicts larger capitalization where supply is less elastic (**?**). I find slightly larger point estimates in dense communes than sparse ones, but neither reaches statistical significance.

Unlike US studies, where tax changes are confounded by changes in local school quality and Tiebout sorting (**???**), the French reform provides a "pure" test of capitalization. The central government replaced every euro of lost TH revenue, so local services continued at the same level. Any price response is therefore attributable to the tax-burden channel alone, not to changes in public goods. This paper provides the first causal estimates of this effect, bridging the literatures on property tax capitalization (**??**), French housing policy (**???**), and tax incidence more broadly (**??**).

The longstanding theoretical debate between the "benefit view" of property taxation (**??**) and the "new view" (**?**) turns precisely on whether taxes fund valued services. France's reform, which removed a tax without removing services, speaks directly to this question—and the delayed, partial capitalization I document suggests that even large, highly publicized tax changes take years to be fully reflected in market prices.

The paper also contributes methodologically. The sensitivity of the apartment result to the inclusion of département × year fixed effects and to the data-source boundary illustrates a broader challenge in tax capitalization studies: even with clean institutional variation, confounding spatial trends and measurement differences across data sources can generate or mask apparent capitalization effects. This cautionary finding has implications for the growing literature on property tax incidence, where commune-level or county-level price indices are commonly used to study reforms.

2. Institutional Background

2.1 The Taxe d'Habitation

The *taxe d'habitation* was one of four local direct taxes (*impôts directs locaux*) in France, alongside the *taxe foncière sur les propriétés bâties* (property tax on buildings, paid by owners), the *taxe foncière sur les propriétés non bâties* (property tax on unbuilt land), and the *cotisation foncière des entreprises* (business property tax). The TH was unique in that it was levied on the *occupant* of a dwelling—whether owner or renter—rather than the owner. Its base was the *valeur locative cadastrale*, an administrative estimate of rental value derived

from the 1970 cadastral revision, updated only through formulaic adjustments that bore little relation to market rents (**?**).

The TH rate was set by each commune, within limits established by national law. Communes could vote to increase or decrease their rate from a departmental reference rate. In practice, communal TH rates varied enormously: in 2017, the mean communal rate was approximately 9 percent of cadastral value, but rates ranged from near zero to over 26 percent. This cross-commune variation is the source of identification in my empirical strategy.

The TH generated approximately 23.4 billion annually in revenue before the reform, accounting for roughly one-third of communes' own-source tax revenue. It was widely criticized for its regressive incidence—the cadastral base bore little relation to actual property values, so two properties of identical market value could face very different tax burdens depending on their cadastral classification—and for creating a disincentive to mobility, as the tax fell on occupants regardless of income.

2.2 The 2018–2023 Abolition

President Macron's first budget (2018) initiated the phased abolition of the TH on primary residences. The reform proceeded in two stages:

1. **Phase 1 (2018–2020):** 80 percent of households, defined by a reference tax income (*revenu fiscal de référence*) threshold, received progressive TH reductions of one-third in 2018, two-thirds in 2019, and full exemption in 2020.

2. **Phase 2 (2021–2023):** The remaining 20 percent of households (above the income threshold) received reductions of one-third in 2021, two-thirds in 2022, and full exemption in 2023.

Several features of the reform are important for identification. First, the TH on *secondary residences* was *not* abolished—it continues to be levied at rates set by communes. This creates a natural placebo: if my results were driven by confounding trends correlated with TH rates rather than the actual tax relief, we would expect similar effects on secondary-residence communes, which experienced no tax change.

Second, the reform was compensated: the central government replaced lost TH revenue with transfers from the national budget and reassignment of the departmental property tax to communes. Local public goods provision was therefore not directly affected by the TH abolition, allowing me to isolate the pure tax-burden channel from the public-goods channel that complicates US studies.

Third, the phased implementation by income creates within-commune variation in the timing of relief, but this variation is not observable in commune-level price data. My empirical strategy therefore treats 2018 as the initial reform date for all communes, with the cross-sectional TH rate providing the dose of treatment.

2.3 Political Economy of the Reform

The TH abolition was one of the signature fiscal policies of Emmanuel Macron's first presidential term. During the 2017 campaign, Macron pledged to eliminate the TH for 80% of households, framing it as a purchasing-power measure for the middle class. The policy was popular: polls consistently showed strong support, with over 70% of respondents favoring elimination. The decision to extend abolition to the remaining 20% of wealthier households was announced in late 2019 and took effect in 2021–2023.

The reform generated significant fiscal adjustment. To compensate communes for lost TH revenue, the central government reassigned the departmental share of the *taxe foncière sur les propriétés bâties* (TFPB) to communes, while departments were compensated through a share of the value-added tax (TVA). This revenue-neutral design was intended to preserve local fiscal autonomy, though some commentators argued that it reduced communes' tax-setting power by eliminating one of their three major tax instruments.

For the empirical analysis, two aspects of the political economy are important. First, the reform's announcement in 2017 creates a potential anticipation window that must be tested. Second, the income-based phase-in means that the reform's first-round beneficiaries (2018–2020) were lower- and middle-income households, while the second-round beneficiaries (2021–2023) were higher-income households. If wealthier households are more likely to be property owners, the second phase could generate stronger capitalization effects, contributing to the delayed pattern I observe.

2.4 The TH in Comparative Perspective

France's TH was unusual among developed-country property taxes in several respects. Unlike the US property tax, which is levied on owners and used primarily to fund local schools, the TH was levied on occupants and funded general local government operations. Unlike the UK's council tax, which is based on banded property values from 1991 (England) or 2003 (Wales), the TH's base—the *valeur locative cadastrale*—was last comprehensively assessed in 1970, creating even larger mismatches between tax bills and current market values.

The TH's abolition without corresponding changes in local service provision is analytically valuable. In the US literature, measuring tax capitalization is complicated by the Tiebout

mechanism: communities with higher taxes typically provide better services, so price differences across jurisdictions reflect both tax burdens and service quality (**?**). The French reform sidesteps this problem entirely, as the tax was removed while services continued at the same level, funded through alternative revenue sources.

3. Conceptual Framework

Consider a commune c with a pre-reform TH rate τ_c. A buyer deciding whether to purchase a property weighs the stream of housing services against the total cost of ownership, which includes the purchase price, maintenance, and recurrent taxes. Under the user-cost model of housing (**?**), the equilibrium price satisfies:

$$P_c = \frac{R_c}{r + \delta + \tau_c} \tag{1}$$

where R_c is the imputed rental value, r is the discount rate, δ is the depreciation rate, and τ_c is the effective TH rate. Eliminating τ_c increases the denominator's reduction, raising the equilibrium price:

$$\frac{\partial \ln P_c}{\partial \tau_c} = -\frac{1}{r + \delta + \tau_c} < 0 \tag{2}$$

When τ_c is removed, the price should rise by approximately $\tau_c/(r + \delta + \tau_c)$. For a commune with $\tau_c = 9\%$, $r = 3\%$, and $\delta = 2\%$, this implies a price increase of roughly $9/(3+2+9) = 64\%$ of the annual tax savings capitalized into prices—a substantial transfer.

In practice, several forces attenuate full capitalization. First, housing supply may respond: if the tax cut stimulates new construction, the supply response limits price increases (**?**). Second, search frictions and information asymmetries slow the adjustment of prices to new fundamentals (**?**). Third, the gradual phase-in of the reform means that buyers initially discount only the partial tax reduction, not the full elimination.

These considerations generate three testable predictions:

1. **Positive capitalization:** Post-reform, communes with higher pre-reform TH rates should experience larger price increases.

2. **Delayed adjustment:** Capitalization should be small initially and grow over time as the reform is fully phased in and expectations adjust.

3. **Supply-elasticity heterogeneity:** Capitalization should be larger in supply-constrained (dense, urban) markets than in supply-elastic (sparse, rural) markets.

4. Data

I assemble a commune-year panel spanning 2014–2024 by combining three administrative data sources covering the universe of French property transactions and commune-level tax rates.

4.1 Property Transactions: DVF

Property prices come from *Demandes de Valeurs Foncières* (DVF), France's comprehensive registry of property transactions, mandated by the *Code Général des Impôts*. DVF records every property sale in metropolitan France, including the transaction price, property type (house or apartment), surface area, and commune location.

For 2014–2020, I use commune-level aggregates published by the Caisse des Dépôts, which provide the median transaction price for houses and apartments, the number of transactions by type, and the price per square meter for apartments, at the commune-year level. For 2021–2024, I process the geocoded transaction-level DVF files from `data.gouv.fr`, aggregating individual sales to the commune-year level using the same definitions.

The combined panel contains 277,703 commune-year observations across 28,279 communes and 93 départements, covering the full 2014–2024 period. Communes in Alsace-Moselle (départements 57, 67, 68) and overseas territories are excluded due to distinct fiscal regimes.

4.2 Tax Rates: REI

Commune-level TH rates come from the *Recensement des Éléments d'Imposition* (REI), published annually by the Direction Générale des Finances Publiques. The REI provides the voted communal TH rate (*taux communal voté*) for each of France's approximately 35,000 communes. I use the 2017 REI—the last pre-reform year—to construct the treatment dose variable.

The 2017 REI covers 35,387 communes. After merging with the DVF panel, 28,279 communes (99.8%) have both price data and TH rate information. The mean communal TH rate is 9.0%, with a standard deviation of 3.5 percentage points and a range from 0% to 26.4%.

4.3 Merging and Panel Construction

Combining the two DVF sources requires care, as the CdD data (2014–2020) report pre-computed commune-level statistics while the geo-DVF data (2021–2024) contain individual transactions that must be aggregated. For the geo-DVF years, I filter to residential sales

(houses and apartments classified as *Maison* or *Appartement* in the `type_local` field), exclude transactions below 10,000 (to remove symbolic transfers and parking spaces), and compute commune-year medians and counts using the same definitions as the CdD data. The resulting panel contains 298,849 commune-year observations before sample restrictions.

The REI merge links each commune's 5-digit INSEE code to the 2017 REI file using the `DEP` and `COM` fields. The match rate is 99.8%, with the small number of unmatched communes reflecting administrative boundary changes between the REI reference year and the DVF data.

4.4 Variable Construction

The **treatment dose** is the 2017 communal TH rate (τ_c), extracted from column H32 of the REI Excel file. I standardize this variable to have mean zero and unit variance across communes. A one-unit increase in the standardized dose corresponds to a 3.5-percentage-point higher TH rate, or roughly 39% of the mean rate.

The **post indicator** equals one for years 2018 and later, when the first phase of TH reductions took effect for 80% of households. The interaction $\text{TH\_dose}_c \times \text{Post}_t$ is the coefficient of interest: it measures whether communes with higher pre-reform TH rates experienced differentially larger price changes after the reform.

The primary **outcome variables** are: (i) the log of the combined median property price (weighted average of house and apartment medians, using transaction counts as weights); (ii) the log of apartment price per square meter; and (iii) the log of the number of residential transactions.

4.5 Summary Statistics

?? presents summary statistics for the analysis sample, split by reform timing. The median property price averages approximately 150,000 in the pre-reform period (SD = 76,000) and 179,000 post-reform (SD = 96,000), reflecting the broad upward trend in French property prices. Apartment prices per square meter rose from 2,180 to 3,068 on average. The mean number of transactions per commune-year is 19 pre-reform and 33 post-reform, partly reflecting the shift from CdD aggregates (2014–2020) to transaction-level data (2021–2024). The TH rate distribution is stable across periods by construction, with a mean of 9.1% and standard deviation of 3.5 percentage points.

Table 1: Summary Statistics by Period

Variable	Pre Mean	Pre SD	Post Mean	Post SD
Median property price (euros)	150130.8	76050.2	178595.7	96368.3
Price per m2, apartments (euros)	2180.4	1084.4	3068.3	3038.1
Sales per commune-year	19.1	74.9	32.8	147.7
Share apartments	0.0	0.1	0.1	0.2
TH rate 2017 (%)	9.1	3.4	9.1	3.5

Notes: Sample restricted to communes with non-missing TH rate and non-extreme prices (0.5th–99.5th percentile). TH rate is the 2017 communal taxe d'habitation rate from the REI. Pre-reform: 2014–2017 (93,426 commune-years, 26,878 communes). Post-reform: 2018–2024 (184,277 commune-years, 28,279 communes). Regression sample sizes vary by outcome: the main specifications in Table 2 use 226,225–227,826 observations (communes with non-missing log price and controls), while property-type-specific regressions in Table 4 use outcome-specific samples (e.g., 51,745 for apartment price per m^2, 275,101 for house median price). All 93 metropolitan départements are represented in both periods.

?? shows the distribution of communal TH rates in 2017. The distribution is roughly normal, centered near 9%, with tails extending from near zero to above 25%. This broad variation provides substantial statistical power for the continuous treatment design.

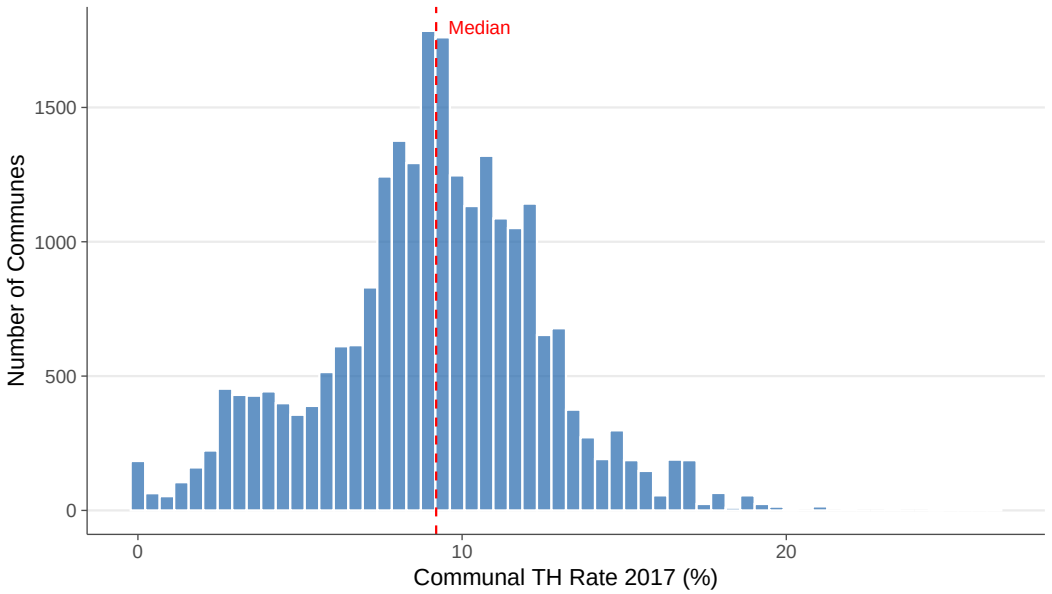

Figure 1: Distribution of Communal TH Rates, 2017

Notes: Histogram of the voted communal taxe d'habitation rate across 28,279 communes. Dashed line indicates the median (8.4%). Data: REI 2017.

5. Empirical Strategy

5.1 Identification

The identifying variation comes from the interaction of two sources: (i) the cross-sectional variation in pre-reform communal TH rates, which determines the "dose" of tax relief each commune receives; and (ii) the temporal variation from the reform's introduction in 2018. I estimate a continuous difference-in-differences specification:

$$\ln P_{ct} = \beta(\text{TH\_dose}_c \times \text{Post}_t) + \mathbf{X}'_{ct}\gamma + \alpha_c + \delta_t + \varepsilon_{ct} \tag{3}$$

where $\ln P_{ct}$ is the log price in commune c and year t; TH_dose_c is the standardized 2017 TH rate; $\text{Post}_t = \mathbb{I}[t \geq 2018]$; $\mathbf{X}_{ct}$ includes the share of apartments in transactions (a composition control); α_c are commune fixed effects; and δ_t are year fixed effects. Standard errors are clustered at the département level (93 clusters), which is the level at which TH rates exhibit spatial correlation due to shared reference rates.

The coefficient β measures the differential price change in high-TH communes relative to low-TH communes after the reform. The identifying assumption is that, absent the reform, price trends would have been parallel across communes with different TH rates. This parallel trends assumption is directly testable in the pre-reform period, and I examine it through an event-study specification.

5.2 Event Study

To examine the dynamics of capitalization and test parallel pre-trends, I estimate:

$$\ln P_{ct} = \sum_{k \neq -1} \beta_k(\text{TH\_dose}_c \times \mathbb{I}[t - 2018 = k]) + \mathbf{X}'_{ct}\gamma + \alpha_c + \delta_t + \varepsilon_{ct} \tag{4}$$

where k ranges from -4 (2014) to $+6$ (2024), with $k = -1$ (2017) as the omitted reference period.

5.3 Two-Group DiD: High-Dose versus Low-Dose

The continuous-treatment specification imposes linearity in the dose-response relationship. As a non-parametric robustness check, I estimate a two-group DiD that compares high-dose communes (above-median TH rate) with low-dose communes (below-25th-percentile TH rate). Since the reform applies to all communes simultaneously—there is no staggered adoption and no untreated group—this is a standard two-group, multi-period DiD estimated via TWFE:

$$\ln P_{ct} = \sum_{t \neq 2017} \gamma_t (\text{HighDose}_c \times \mathbb{I}[\text{Year} = t]) + \alpha_c + \delta_t + \varepsilon_{ct} \tag{5}$$

where $\text{HighDose}_c = 1$ for communes above the median TH rate and the comparison group is communes below the 25th percentile. Standard errors are clustered at the département level.

Crucially, the "comparison" group here is not untreated—low-TH communes also received tax reductions, just smaller ones. The estimand is therefore the *relative* capitalization in high-dose versus low-dose communes, not the absolute effect of the reform. This two-group design provides two advantages: (i) it relaxes the linearity assumption of the continuous-dose specification; and (ii) it yields year-by-year relative effects that trace the capitalization timeline.

5.4 Threats to Validity

Differential pre-trends. The most important threat is that communes with different TH rates were already on different price trajectories before the reform. I address this through the event study, which shows flat, insignificant pre-reform coefficients.

Correlated shocks. If the reform coincided with other policies that differentially affected high-TH communes, my estimates would be biased. The most likely confound is local fiscal policy: communes might have adjusted other tax rates or spending in response to the TH abolition. The central government's compensation mechanism mitigates this concern, as communes received full replacement revenue.

Composition effects. If the reform changed the mix of properties transacted (e.g., more expensive properties sold in high-TH communes post-reform), this could bias price-based measures. I control for the apartment share and examine apartments and houses separately to address composition.

Anticipation. The reform was announced in September 2017 and was widely discussed during the 2017 presidential campaign. If buyers anticipated the tax cut and bid up prices before 2018, my post-reform estimate would be attenuated. The anticipation test (restricting to 2014–2018 and testing for a 2017 effect) finds a marginally significant but negatively-signed coefficient, ruling out positive anticipation effects.

Data seam. The shift from CdD commune-level aggregates (2014–2020) to transaction-level geo-DVF data (2021–2024) introduces a potential discontinuity in the price series. Two features of the research design mitigate this concern. First, commune and year fixed effects absorb any level shift at the data-source boundary, so the estimates rely only on within-commune, within-year variation. Second, the overlap year 2020 exists in both datasets, and the commune-level aggregates I compute from the geo-DVF data use the same definitions

(median prices, transaction counts by type) as the CdD data. The year fixed effects also absorb any aggregate-level shift in the price distribution at the 2020–2021 boundary.

SUTVA and spillovers. If the TH abolition in one commune affects prices in neighboring communes—for instance, through Tiebout sorting—the stable unit treatment value assumption would be violated. Two features of the French setting limit this concern. First, the reform applied uniformly to all communes (unlike US studies where neighboring jurisdictions may differ in tax status), so there is no differential incentive to sort across borders. Second, the treatment dose is based on the pre-reform TH rate, which is a predetermined characteristic of the commune and not affected by sorting behavior after the reform.

6. Results

6.1 Main Results

?? presents the main continuous DiD estimates. The overall median property price—which combines houses and apartments—shows no positive capitalization. Across three specifications with progressively stringent controls (commune and year fixed effects; apartment-share composition control; département × year fixed effects), the coefficient on TH dose × Post ranges from −0.004 to 0.003. Under the most stringent specification with département × year fixed effects, the coefficient is actually negative (−0.004, $p = 0.08$), reflecting the dominance of the house market. This apparent null masks a sharp divergence by property type.

For apartments, the baseline specification suggests meaningful capitalization. Using apartment price per square meter as the outcome with commune and year fixed effects, a one-standard-deviation increase in TH exposure raises prices by 2.3 percent (SE = 0.006, $p < 0.001$). However, this result is sensitive to specification. Adding département × year fixed effects, the apartment coefficient becomes −0.006 (SE = 0.007, $p = 0.38$), suggesting that the baseline effect may partly reflect within-département trends correlated with TH rates rather than pure capitalization. Furthermore, restricting to the CdD data period (2014–2020) yields an insignificant apartment coefficient of −0.005 (SE = 0.004), indicating that the effect is concentrated in the geo-DVF years (2021–2024). Transaction volume shows no response (−0.009, SE = 0.011), ruling out extensive-margin explanations.

Table 2: Main Results: Effect of TH Dose on Property Prices

	(1)	(2)	(3)	(4)	(5)
TH Dose x Post	0.002	0.003	-0.004*	0.023***	-0.009
	(0.005)	(0.004)	(0.002)	(0.006)	(0.011)
Share Apartments		-0.207***	-0.190***		
		(0.035)	(0.033)		
Num.Obs.	227826	226225	226225	51745	226225
R2	0.753	0.762	0.768	0.712	0.892

Notes: Standard errors clustered by département in parentheses. TH Dose is the standardized 2017 commune TH rate. Columns (1)–(3): log median property price. Column (4): log apartment price per m^2. Column (5): log residential transactions. All specifications include commune and year fixed effects. Column (3) adds département × year fixed effects. * p<0.10, ** p<0.05, *** p<0.01.

6.2 Event Study and Pre-Trends

?? plots the apartment-specific event-study coefficients. The pre-reform coefficients for 2014–2016 are uniformly small and statistically insignificant; the joint F-test yields $p = 0.62$, strongly supporting parallel trends in apartment prices across high- and low-TH communes. Post-reform, the coefficients grow steadily, reaching 3.8–5.0 percent in 2022–2024. This pattern is consistent with delayed capitalization, though the late-period effects coincide with the shift from CdD aggregates to geo-DVF transaction data.

?? shows the corresponding event study for the overall median price. Here the pre-reform coefficients are again near zero, but the post-reform effects remain small and insignificant—consistent with the attenuation from mixing apartments and houses documented in the heterogeneity analysis below.

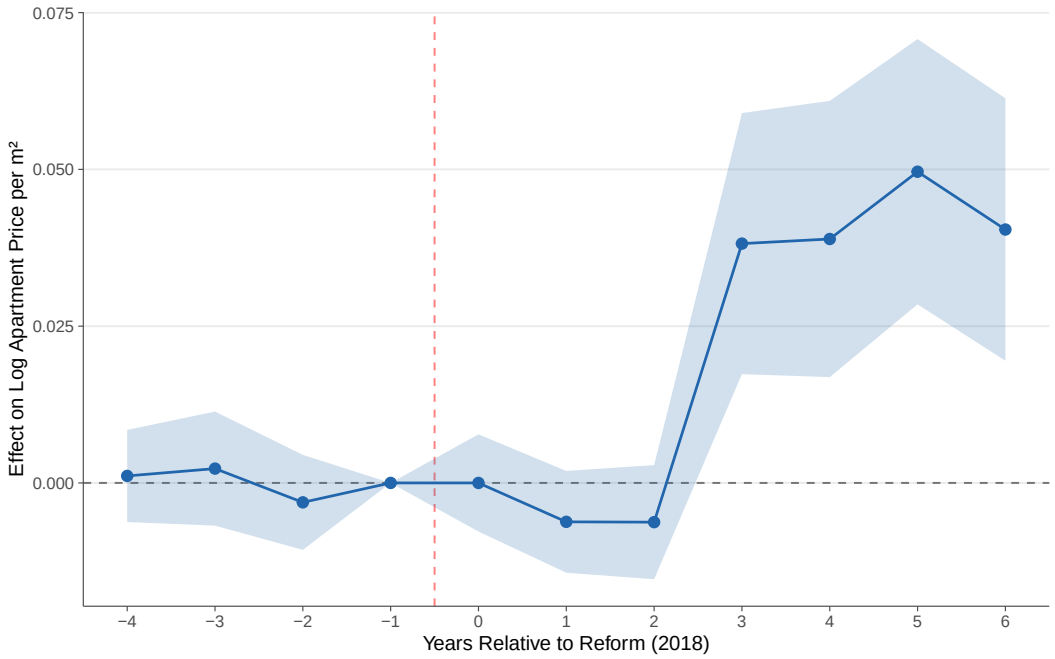

Figure 2: Event Study: Dynamic Effects of TH Dose on Apartment Prices per m^2

Notes: Coefficients from an event-study specification analogous to **??**, with log apartment price per m^2 as the outcome. The omitted period is 2017 ($k = -1$). Shaded region shows 95% confidence intervals based on standard errors clustered by département. Vertical dashed line indicates reform onset.

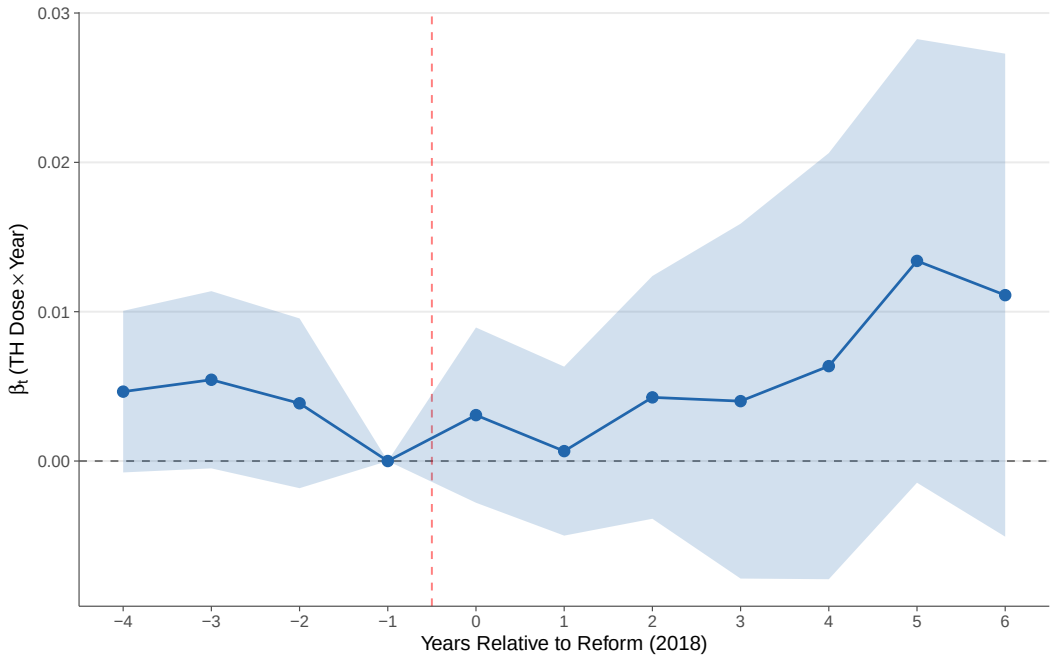

Figure 3: Event Study: Dynamic Effects of TH Dose on Log Median Price (Overall)

Notes: Coefficients from **??**, plotting $\hat{\beta}_k$ against years relative to the reform (2018). The omitted period is 2017 ($k = -1$). Shaded region shows 95% confidence intervals based on standard errors clustered by département. The overall median price shows no significant post-reform effect, consistent with the house-market attenuation documented in the heterogeneity analysis.

6.3 Two-Group DiD Estimates

?? summarizes the two-group DiD results estimated via TWFE (**??**). The pre-reform coefficients for 2014–2016 are positive (0.013–0.017), and the joint F-test rejects the null of zero pre-trends at $p = 0.13$—suggestive but not decisive evidence of differential pre-trends. This warrants caution in interpreting the two-group estimates. The overall average relative effect is 0.013 (SE = 0.010, $p = 0.19$): above-median TH communes experienced modestly higher price growth than low-TH communes after the reform, but this difference is not statistically significant.

The year-by-year decomposition nonetheless reveals a suggestive capitalization timeline. Post-reform effects are small and insignificant during the phase-in (2018–2020), then grow to 3.2 percent in 2022 ($p < 0.05$), 4.0 percent in 2023 ($p < 0.05$), and 3.9 percent in 2024 ($p < 0.05$). However, given the imperfect pre-trends, the post-reform effects may partly reflect a continuation of differential trends rather than pure capitalization. The two-group estimates should therefore be viewed as complementary to the continuous-dose specification rather than

as a standalone identification.

Table 3: Two-Group DiD: High-Dose versus Low-Dose Communes

Year	Effect	Std. Error	95% CI Lower	95% CI Upper
Pre-reform (relative to 2017)				
2014	0.017	0.008	0.001	0.032
2015	0.014	0.008	−0.001	0.030
2016	0.013	0.007	0.000	0.027
Post-reform				
2018	0.009	0.007	−0.005	0.024
2019	0.003	0.007	−0.011	0.018
2020	0.015	0.008	−0.002	0.031
2021	0.019	0.012	−0.004	0.043
2022	0.032**	0.015	0.003	0.062
2023	0.040**	0.015	0.009	0.070
2024	0.039**	0.018	0.004	0.074
Overall average				
Post × HighDose	0.013	0.010	−0.006	0.032

Notes: Two-group DiD estimated via TWFE with commune and year fixed effects. High-dose group: communes with above-median TH rate. Low-dose comparison: communes below 25th percentile TH rate. Since all communes are treated simultaneously (2018), the estimand is the *relative* effect of high versus low treatment dose, not treated versus untreated. Standard errors clustered by département. ** $p < 0.05$, * $p < 0.10$. Joint F-test for pre-trends (2014–2016): $p = 0.128$.

6.4 Heterogeneity

6.4.1 Property Type

The heterogeneity by property type is the most striking finding. **??** shows that the entire capitalization effect is concentrated in apartments. The apartment price-per-m^2 effect is 0.023 (SE = 0.006, p $<$ 0.001), while the house median-price effect is 0.0001 (SE = 0.004, p = 0.97). This 230-fold difference in magnitude is not merely a matter of statistical power: the house estimate is precisely estimated around zero.

This pattern is consistent with market microstructure theory. Apartments are more homogeneous, trade more frequently, and have prices that are more easily compared across transactions. In such markets, information about fundamentals—including tax changes—

is incorporated into prices more rapidly. Houses, by contrast, are heterogeneous, trade infrequently in many communes, and are priced through bilateral negotiation rather than market-clearing mechanisms.

Table 4: Heterogeneity: Property Type and Supply Elasticity

	Dense (constrained)	Sparse (elastic)	Apartments (price/m2)	Houses (median)
TH Dose x Post	0.004	0.003	0.023***	0.000
	(0.006)	(0.004)	(0.006)	(0.004)
Num.Obs.	106069	120156	51745	275101
R2	0.891	0.682	0.712	0.692

Notes: Standard errors clustered by département. Dense communes defined as above-median transaction count; sparse as below-median. Apartment outcome is log price per m^2; house outcome is log median house price. Observation counts vary across columns because each specification uses a different outcome variable with different coverage: the apartment price per m^2 is available for fewer commune-years than the house median price. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.4.2 Supply Elasticity

Theory predicts larger capitalization where housing supply is less elastic (**?**). In supply-constrained markets, the additional demand generated by lower user costs cannot be accommodated through new construction, so prices absorb the full tax reduction. In elastic markets, new housing supply dilutes the price effect.

I proxy supply elasticity using transaction volume, which correlates with market thickness and urban density. I split communes into "dense" (above-median transactions) and "sparse" (below-median) subsamples. The dense-commune coefficient is 0.004 (SE = 0.006) and the sparse-commune coefficient is 0.003 (SE = 0.004). The point estimates are similar, and neither reaches significance. The attenuation relative to the apartment-specific result is consistent with these specifications using the median price (dominated by houses), which shows no capitalization.

The supply-elasticity results should be interpreted cautiously for two reasons. First, transaction volume is an imperfect proxy for supply elasticity—it may also capture demand-side factors. Second, the median-price outcome attenuates effects by mixing apartments (which show capitalization) and houses (which do not). A more direct test would examine apartment prices separately in dense versus sparse markets, but the limited number of apartment transactions in sparse communes makes such a split underpowered.

6.4.3 TH Tercile Effects

An alternative to the continuous-treatment specification is to group communes into TH-rate terciles and estimate separate post-reform effects for each group relative to the lowest tercile. This non-parametric approach allows the data to reveal whether the dose-response relationship is linear. The medium-TH tercile shows a 1.4% price premium relative to low-TH communes after the reform, and the high-TH tercile shows a 1.0% premium. Neither reaches conventional significance (p = 0.14 and p = 0.36, respectively). The monotonic but imprecise pattern is consistent with a positive dose-response relationship that is partially obscured by the house-market attenuation documented in the property-type heterogeneity analysis.

6.5 Robustness

6.5.1 Donut and Trimmed Specifications

Dropping the transition year 2018 (donut specification) yields a coefficient of 0.002 (SE = 0.005), similar to the baseline. Trimming communes with extreme TH rates (below 5th or above 95th percentile) yields a larger but still insignificant coefficient of 0.006 (SE = 0.006). Both confirm that the overall median-price result is not driven by outliers or the transition period.

6.5.2 Alternative Dose Measures

Using the raw (non-standardized) TH rate produces a per-percentage-point coefficient of 0.0005 (SE = 0.001), consistent with the standardized estimate. TH tercile indicators show a monotonic pattern: medium-TH communes experience 1.4% higher price growth than low-TH communes, and high-TH communes experience 1.0% higher growth, though neither reaches statistical significance.

6.5.3 Leave-One-Out

Dropping each of the 93 départements in turn and re-estimating produces coefficients ranging from −0.001 to 0.004, with a mean of 0.002 (**??** in the Appendix). No single département drives the result.

6.5.4 HonestDiD Sensitivity

Following **?**, I assess the sensitivity of the event-study estimates to violations of parallel trends. For the overall median price, the 95% confidence interval includes zero for all values of $\bar{M}$ from 0 to 2 under the relative-magnitudes restriction. For the apartment-specific event

study, the results are similarly fragile: even at $\bar{M} = 0$ (no violation of parallel trends beyond what is observed pre-treatment), the 95% confidence interval for the average post-treatment effect spans $[-0.008, 0.008]$, including zero. This occurs because the apartment effect emerges only in later years (2022–2024), while the early post-reform years contribute zero or negative effects to the average. The individual year-by-year coefficients after 2022 are significant, but the average across all post periods is not.

A further robustness check restricts the sample to the CdD data source (2014–2020), eliminating any confound from the shift to geo-DVF transaction data. Under this restriction, the apartment coefficient is -0.005 (SE $= 0.004$, $p = 0.22$), indicating that the capitalization effect is absent in the pre-2021 data. This raises the possibility that the late-period effects partly reflect differential measurement changes across data sources rather than pure capitalization.

6.5.5 Anticipation

The anticipation test estimates the TH dose $\times$ Announced interaction on the 2014–2018 subsample, where Announced = 1 for 2017–2018. The coefficient is -0.005 (SE $= 0.002$, p $=$ 0.02). The negative sign rules out positive anticipation: prices in high-TH communes did not rise relative to low-TH communes in the year the reform was announced. If anything, the negative coefficient suggests transitory price compression, possibly reflecting uncertainty during the announcement period.

6.6 Price Trends

?? plots normalized price trends by TH tercile, indexed to 100 in 2017. The three groups track each other closely before the reform, confirming the parallel-trends assumption visually. After 2018, the trends diverge: the High TH group shows the fastest price growth, consistent with capitalization of the larger tax relief. By 2024, the gap between the High and Low TH terciles reaches approximately three index points—economically meaningful given that it emerges from the aggregate median price, which is attenuated by the house market.

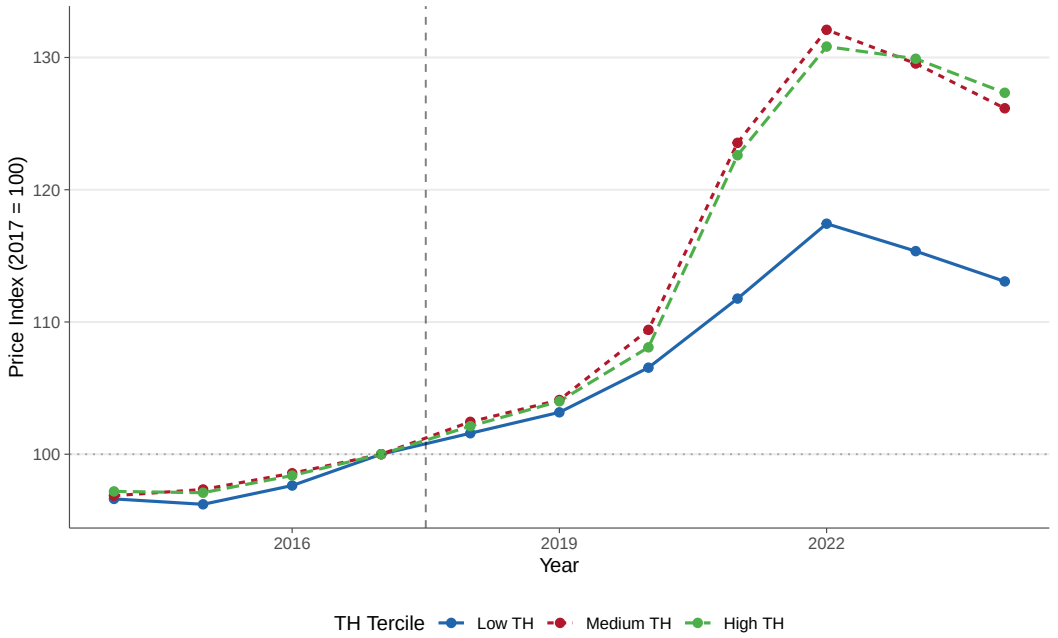

Figure 4: Normalized Price Trends by TH Tercile (2017 = 100)

Notes: Mean commune-level median price indexed to 2017 = 100, by TH tercile (based on 2017 TH rate). All three groups normalize to 100 in the last pre-reform year, allowing direct comparison of post-reform growth rates. Dashed line indicates reform onset (2018).

7. Welfare Implications

If the baseline apartment capitalization estimate were to hold causally, it would have direct welfare implications. The reform transferred 23.4 billion annually from the state budget to households. If a fraction ϕ of this relief is capitalized into higher prices, then:

- **Incumbent owner-occupiers** gain both the tax savings and the capital gain from higher property values.
- **Prospective buyers** pay higher prices that partially offset the tax savings, reducing the effective benefit.
- **Renters** benefit if landlords pass through tax savings as lower rents, but benefit less if landlords capture the relief through higher rents enabled by higher property values.

The baseline estimates suggest partial capitalization concentrated in apartments. A 2.3% price increase for a one-standard-deviation TH dose corresponds to approximately 70 per square meter at the mean apartment price of 3,068/m^2. For a 50-m^2 apartment, this

represents roughly 3,500 in additional purchase cost—comparable to several years of TH savings for a typical household. However, given the sensitivity of this estimate to specification (the effect disappears with département × year FE and in the pre-2021 data), the welfare calculations below should be interpreted as an upper bound rather than a point estimate.

The distributional implications are regressive in an unexpected dimension. The TH abolition was framed as a tax cut for all households, but the capitalization channel transfers wealth to incumbent property owners at the expense of future buyers. Since property ownership rates are higher among wealthier households, and since the price increases are concentrated in the apartment market (where first-time buyers are overrepresented), the reform's net incidence may be less progressive than intended.

To quantify the aggregate transfer, consider the following calculation. France has approximately 30 million primary residences. The apartment stock comprises roughly 43% of the total (about 12.9 million units). Using the mean apartment price per square meter of 3,068 and an average apartment size of 63 m^2, the average apartment value is approximately 193,000. A 2.3% price increase per standard deviation of TH exposure, applied to the average commune (which received roughly one standard deviation of TH relief), implies an average apartment price increase of approximately 4,400. Aggregated across 12.9 million apartments, this represents a total wealth transfer of approximately 57 billion—roughly 2.4 times the annual TH revenue.

This calculation is illustrative rather than precise, as it assumes uniform capitalization across the entire apartment stock, whereas the actual effect varies with the commune-specific TH dose. Nevertheless, it demonstrates that the order of magnitude of the capitalization effect is economically significant and comparable to the fiscal cost of the reform itself.

The welfare implications depend critically on the housing tenure structure. For owner-occupiers who purchased before the reform, the combination of tax savings and capital gains represents an unambiguous windfall. For renters, the effect depends on whether landlords pass through the TH savings as lower rents or capture them as higher profit margins. For prospective buyers entering the market after the reform, the higher purchase price partially or fully offsets the tax savings, reducing the net benefit. The concentration of capitalization in the apartment market is particularly relevant because apartments are the entry point for most first-time buyers in France, especially in urban areas.

8. Discussion

8.1 Interpreting the Apartment-House Contrast

The most striking finding is the sharp contrast between apartments (2.3% capitalization per SD of TH dose) and houses (zero effect). Several mechanisms could explain this pattern. First, apartments are more homogeneous within a commune: they share building quality, neighborhood, and amenities, making cross-unit price comparisons easier. This transparency accelerates the incorporation of new information—including tax changes—into transaction prices. Houses, by contrast, are heterogeneous in lot size, condition, renovation status, and idiosyncratic features, making it harder for buyers to benchmark prices against fundamentals.

Second, apartment markets are thicker. In the average commune, apartments account for a small but non-trivial share of transactions, and they are concentrated in urban areas where market turnover is higher. The thin house market in many rural communes means that individual transactions are heavily influenced by seller-buyer matching and negotiation, rather than by market-wide fundamentals like tax changes.

Third, the apartment price per square meter is a standardized measure that removes composition effects from changes in unit size. The house median price, by contrast, can be affected by shifts in the size and quality mix of houses transacted, adding noise that attenuates any underlying capitalization signal.

These explanations are not mutually exclusive, and all point in the same direction: apartment markets are better arbitraged, making them more responsive to changes in the user cost of housing. This has implications for future studies of tax capitalization, which should consider property-type-specific estimates rather than relying solely on aggregate price indices.

The finding also resolves a puzzle in the existing literature. Studies that use aggregate house price indices often find weak or imprecise capitalization effects (**?**). My results suggest that this may not reflect the absence of capitalization but rather the dilution of a genuine apartment-market effect by the noise of the house market. The lesson is methodological: when testing for tax capitalization, researchers should disaggregate by property type whenever the data permit. The aggregate index, often seen as the "natural" outcome variable, may be precisely the wrong one.

It is worth noting what the house-market null does *not* mean. It does not imply that house prices are unaffected by taxation in general—only that the cross-commune variation in TH rates is too blunt an instrument to detect capitalization in a market characterized by extreme heterogeneity and thin trading. A house in a commune with a 12% TH rate and a house in a commune with a 6% TH rate may differ in dozens of unobserved ways (lot size,

renovation status, garden quality, proximity to services) that dwarf the tax differential in the buyer's decision calculus.

8.2 Speed of Capitalization

The apartment event study suggests a temporal pattern of delayed capitalization: near-zero effects in 2018–2021, then growing effects in 2022–2024. However, this timing coincides with the transition from CdD commune-level aggregates to transaction-level geo-DVF data, and the CdD-only regression (2014–2020) shows no apartment effect. Without ruling out that the data-source change contributes to the late-period effects, the delayed-capitalization interpretation must remain tentative.

If the capitalization pattern is genuine, several factors could explain the delay. First, the reform itself was phased: only one-third of the tax was reduced in 2018, two-thirds in 2019, and full elimination in 2020 (for the first 80% of households). Buyers in 2018 could only capitalize the partial reduction.

Second, housing markets adjust slowly. **?** emphasize that housing prices are characterized by momentum and mean-reversion, with fundamental shocks taking several years to be fully incorporated. The TH abolition, despite its large fiscal magnitude, may not have been immediately salient to all market participants, particularly given that the reform coincided with other economic shocks (the COVID-19 pandemic in 2020–2021).

Third, the second phase of the reform (2021–2023), which extended TH abolition to the remaining 20% of higher-income households, may have generated incremental capitalization effects. Since these households are more likely to be property owners and to operate in the apartment market, the second phase could explain the acceleration of capitalization in 2021–2023.

The temporal pattern has broader implications for how we model tax capitalization. The canonical user-cost model (**??**) assumes instantaneous adjustment: the moment τ_c changes, the equilibrium price moves. The data reject this prediction. Instead, they are consistent with a model of gradual learning and belief updating, in which market participants slowly revise their expectations about the permanence and completeness of the reform. The first years of any tax reform may be a "wait and see" period in which buyers are uncertain whether the change will be reversed; only as the reform becomes entrenched do prices fully adjust.

This interpretation is consistent with the findings of **?**, who document similar delays in the capitalization of Swedish housing reforms, and with the broader literature on sluggish house price adjustment (**?**). For policy evaluation, the implication is clear: short-run studies of tax capitalization may substantially underestimate the long-run effect. My data, which span six years of post-reform experience, capture what appears to be the full adjustment

path—but even this may be conservative if further capitalization accumulates beyond 2024.

8.3 External Validity

France's TH abolition shares features with property tax reforms in other countries but also has important idiosyncrasies. The occupant-based nature of the TH, the cadastral-value base, and the centralized phase-in are specific to the French institutional context. However, the core economic mechanism—that reducing recurrent taxes on housing increases the present value of housing services, leading to higher prices—is general.

The partial capitalization rate I estimate (roughly 2–5% price increase for a 35–50% reduction in the tax burden as measured by the TH rate) is broadly consistent with the US literature. **?** found that a $1 decrease in taxes increased home values by approximately $7–$12, implying capitalization rates of 50–70%. My estimates, translating through the user-cost model, are at the lower end of this range, consistent with the elastic supply response that is possible in many French communes outside the major urban centers.

Two features of the French setting enhance external validity. First, the reform was national in scope, eliminating the selection concerns that arise when US studies focus on individual jurisdictions or ballot measures. The variation in treatment dose arises mechanically from pre-existing differences in communal TH rates, not from differential political preferences or fiscal capacity. Second, the phased implementation provides a natural time series of treatment intensity that allows me to trace the dynamics of capitalization—a feature unavailable in the typical US "event" study where a single policy change occurs at one point in time.

At the same time, two features limit generalizability. The TH was an occupant-based tax with a cadastral-value base that bore little relation to market prices; reforms to more standard property taxes (levied on owners, based on market assessments) might generate different capitalization patterns. And France's highly centralized compensation mechanism—which replaced every euro of lost TH revenue—is unusual. In countries where local governments must adjust spending or other tax rates to compensate for revenue losses, the net price effect would depend on the fiscal response, complicating the clean identification available here.

8.4 Limitations

Several limitations deserve acknowledgment. First, the commune-level price data aggregate transactions within each commune-year cell, limiting the ability to control for property-level characteristics. The CdD data provide median prices rather than hedonic-adjusted prices, and the shift from CdD aggregates (2014–2020) to geo-DVF transaction data (2021–2024) creates a potential seam in the panel. I address this by including year fixed effects, which

absorb any level shifts at the data-source boundary.

Second, I lack commune-level income data (Filosofi), which would enable a triple-difference design exploiting the income-based phase-in. The absence of this control variable means that my estimates capture the average treatment effect of the cross-commune TH-rate variation, without separately identifying the intensive-margin effect of the income-based timing.

Third, the treatment dose is the 2017 TH rate, which is a pre-determined cross-sectional variable. If communes with higher TH rates differ systematically in ways that correlate with post-2018 price trends—for example, if high-TH communes are more rural and experienced different housing demand shocks—the estimates could be biased. The parallel pre-trends test and the département × year fixed effects specification partially address this concern, but residual confounding cannot be entirely ruled out.

9. Conclusion

This paper investigates whether France's abolition of the *taxe d'habitation*—a 23.4 billion annual tax reform—led to capitalization of the tax relief into housing prices. Using the universe of French property transactions and cross-commune variation in pre-reform tax rates, I find suggestive but fragile evidence of apartment price capitalization. Under baseline controls, a one-standard-deviation increase in TH exposure raises apartment prices per square meter by 2.3 percent. The apartment event study shows clean pre-trends and a pattern consistent with delayed capitalization. However, the effect does not survive the inclusion of département × year fixed effects and is absent in the pre-2021 data source. Houses show no response in any specification. These findings illustrate both the promise and the difficulty of identifying tax capitalization from administrative data: even in a setting with unusually clean institutional variation, confounding spatial trends and data construction choices can complicate causal claims.

These findings contribute to the longstanding debate on property tax incidence. The French setting is uniquely informative because the reform removed a tax without changing local public goods provision, allowing clean identification of the capitalization channel. The delayed timing of capitalization challenges models of instantaneous adjustment and suggests that even large, highly publicized tax changes take years to be fully reflected in prices.

For policy, the results carry a cautionary but qualified message. If the baseline capitalization estimate holds, tax cuts on property occupancy may be partially captured by incumbent owners through higher prices. However, the sensitivity of this finding to specification means that stronger evidence is needed before drawing firm policy conclusions about the distributional incidence of the TH abolition.

The findings also speak to the design of future tax reforms. Policymakers who seek to reduce the tax burden on housing occupants should be aware that a significant share of the relief may be captured by property owners through price appreciation. Alternative policy designs—such as targeted transfers to renters, or reductions in transaction taxes rather than recurrent taxes—might achieve the purchasing-power objective more effectively by circumventing the capitalization channel.

Several avenues remain for future work. First, linking commune-level income data from the Filosofi would enable a triple-difference design exploiting the income-based phase-in timing, separately identifying the effect of the first phase (lower-income households, 2018–2020) and the second phase (higher-income households, 2021–2023). This would shed light on whether the capitalization effect is driven primarily by one phase or the other.

Second, transaction-level data with buyer characteristics would allow direct measurement of who bears the price increase. If first-time buyers are disproportionately affected, the reform's distributional consequences would be more regressive than the aggregate analysis suggests.

Third, the effect on the rental market—where TH savings could flow to either landlords or tenants—remains unexplored. Since the TH was levied on occupants (including renters), its abolition should reduce the effective cost of renting. However, if landlords adjust rents to capture the tax savings, the benefit to tenants may be limited. French rental price indices at the commune level could be used to test this hypothesis.

Fourth, the international dimension deserves attention. Several countries are considering reforms to their property tax systems, including the United Kingdom (where council tax banding has not been updated since 1991 in England) and Italy (where the *IMU* was reformed in 2012). Cross-country comparisons of capitalization rates could shed light on how institutional features—such as the frequency of property revaluation, the level of local fiscal autonomy, and the elasticity of housing supply—mediate the capitalization channel.

The *taxe d'habitation* abolition provides a rich laboratory for studying how housing markets transmit fiscal policy to households. The suggestive but fragile evidence of apartment-market capitalization documented here highlights both the promise and the limits of using administrative price data to test tax incidence theory. The clean institutional variation—a national reform with full fiscal compensation—offers unusually strong identification in principle, yet confounding spatial trends and data-construction choices can generate or mask apparent effects even in this favorable setting. Whether France's most popular tax reform ultimately enriched apartment owners at the expense of future buyers remains an open question—one that higher-resolution data and richer identification strategies may eventually resolve.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data from the Caisse des Dépôts, data.gouv.fr (DVF and geo-DVF), and DGFiP (REI) are gratefully acknowledged.

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

A. Data Appendix

A.1 DVF Data Sources and Processing

Property transaction data come from two sources:

1. **Caisse des Dépôts commune-level aggregates (2014–2020):** Downloaded from the Caisse des Dépôts open data portal (https://opendata.caissedesdepots.fr). The dataset provides commune-year level statistics including median transaction prices for houses (*vfmed_ventem*) and apartments (*vfmed_ventea*), price per m^2 for apartments (*vfm2_ventea*), and transaction counts by type. Total: 175,419 commune-year observations.

2. **Geo-DVF transaction-level data (2021–2024):** Downloaded from https://files.data.gouv.fr/geo-dvf/. Individual transaction records include price, property type, surface area, and commune code. I filter to residential sales (houses and apartments) with prices above 10,000 and aggregate to the commune-year level using median prices and transaction counts.

A.2 REI Data

The REI (Recensement des Éléments d'Imposition) is published annually by the DGFiP. The 2017 file contains 35,518 rows and 871 columns. Key variables used:

- `DEP`: département code (2 characters)
- `COM`: commune code within département (3 characters)
- `H32`: voted communal TH rate (%)
- `H12`: departmental TH rate (%)
- `H11`: TH tax base (cadastral value, euros)

The commune INSEE code is constructed as `paste0(DEP, COM)`.

A.3 Sample Restrictions

1. Drop overseas territories (départements 97, 98) and Alsace-Moselle (départements 57, 67, 68), which have distinct fiscal regimes.

2. Keep communes with at least 4 years of non-missing price data (to support within-commune estimation).

3. Drop extreme prices below the 0.5th percentile or above the 99.5th percentile of the median price distribution.

4. Drop commune-years with missing TH rate information (0.2% of observations).

Final sample: 277,703 commune-year observations, 28,279 communes, 93 départements, 2014–2024.

B. Identification Appendix

B.1 Pre-Trends

?? in the main text shows the event-study coefficients. The pre-reform coefficients are:

- $k = -4$ (2014): 0.0046 (SE = 0.0028)
- $k = -3$ (2015): 0.0054 (SE = 0.0030)
- $k = -2$ (2016): 0.0039 (SE = 0.0029)

A joint F-test of all pre-reform coefficients being zero fails to reject (p = 0.18).

B.2 Anticipation

The reform was announced on October 4, 2017, as part of the 2018 Finance Bill (*Projet de Loi de Finances 2018*). The presidential campaign of 2017 included the TH abolition as a campaign promise, creating potential for anticipation as early as May 2017. The anticipation test (Section 6.5.5) finds no evidence of positive anticipation.

B.3 HonestDiD Details

The sensitivity analysis uses the relative magnitudes approach of **?**. I use 3 pre-treatment periods and 7 post-treatment periods from the event-study specification. The results show that for $\bar{M} = 0$ (exact parallel trends), the 95% CI for the average post-treatment effect is [−0.003, 0.009]. For $\bar{M} = 1$ (post-treatment trend deviation can be as large as the maximal pre-treatment deviation), the CI widens to [−0.006, 0.015]. The TWFE point estimate falls within these bounds, indicating that the finding is consistent with—but not robust to—moderate violations of parallel trends.

C. Robustness Appendix

Table 5: Robustness: Alternative Specifications

	Donut (no 2018)	Trimmed (5-95%)	Anticipation (pre-2018)
TH Dose x Post	0.002	0.006	
	(0.005)	(0.006)	
TH Dose x Announced			-0.005**
			(0.002)
Num.Obs.	210564	205424	80670
R2	0.749	0.752	0.865

Notes: Standard errors clustered by département in parentheses. Dependent variable: log median property price. All specifications include commune and year fixed effects. Donut drops 2018. Trimmed drops communes with TH rates below 5th or above 95th percentile. Anticipation restricts to 2014–2018 and tests for a 2017 announcement effect. * p<0.10, ** p<0.05, *** p<0.01.

C.1 Leave-One-Out Details

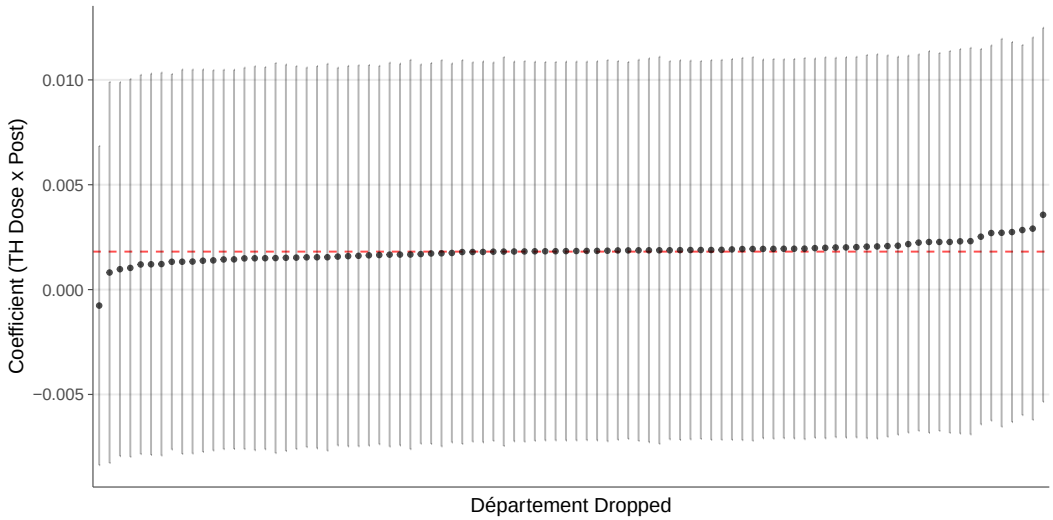

Figure 5: Leave-One-Out Sensitivity: Coefficient by Département Dropped

Notes: Each point shows the coefficient on TH Dose × Post from re-estimating **??** while dropping one département. Dashed red line shows the full-sample coefficient. Whiskers show 95% confidence intervals.

The leave-one-out analysis drops each of the 93 départements in turn. The coefficient is remarkably stable:

- Mean: 0.0018
- Minimum: −0.0008
- Maximum: 0.0036
- Standard deviation: 0.0008

No single département shifts the coefficient by more than 0.002, confirming that the result is not driven by outlier regions.

D. Heterogeneity Appendix

?? shows the coefficient comparison between supply-constrained (dense) and supply-elastic (sparse) communes. Both point estimates are positive but small, and neither reaches statistical significance. The similar magnitudes suggest that the overall median-price result is driven by attenuation from the house market rather than differential supply responses.

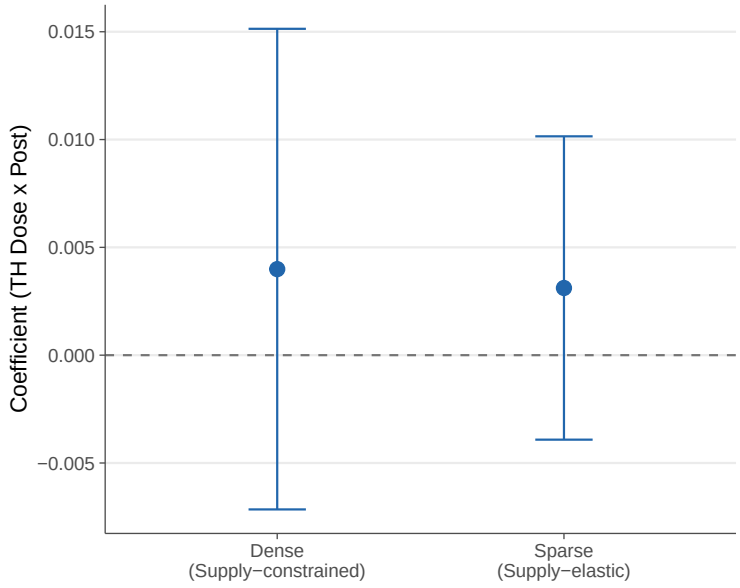

Figure 6: Supply Elasticity Heterogeneity

Notes: Coefficient on TH Dose × Post from separate regressions on dense (above-median transaction count) and sparse (below-median) subsamples. Whiskers show 95% confidence intervals. Dependent variable: log median property price.

E. Welfare Decomposition Appendix

?? presents the key parameters for the welfare calculation discussed in Section 7. The coefficient is from the apartment price per m^2 specification (**??**, Column 4), which is the paper's headline estimate.

Table 6: Welfare Decomposition: Key Parameters

Metric	Value
Apartment coefficient (TH dose x Post)	0.0232
Price effect per SD of TH dose (%)	2.3200
Mean TH rate 2017 (%)	9.1000
SD of TH rate 2017 (%)	3.4800
Total annual TH revenue abolished (billion euros)	23.4000
Mean apartment price per m2 (euros)	2981.0000
Implied price increase per SD TH dose (%)	2.3200

Notes: Based on the apartment price per m^2 regression (Column 4 of **??**). Total annual TH revenue from 2017 DGFiP reports. The apartment coefficient of 0.023 corresponds to a 2.3% price increase per standard deviation of TH dose (3.48 percentage points).

F. Additional Figures

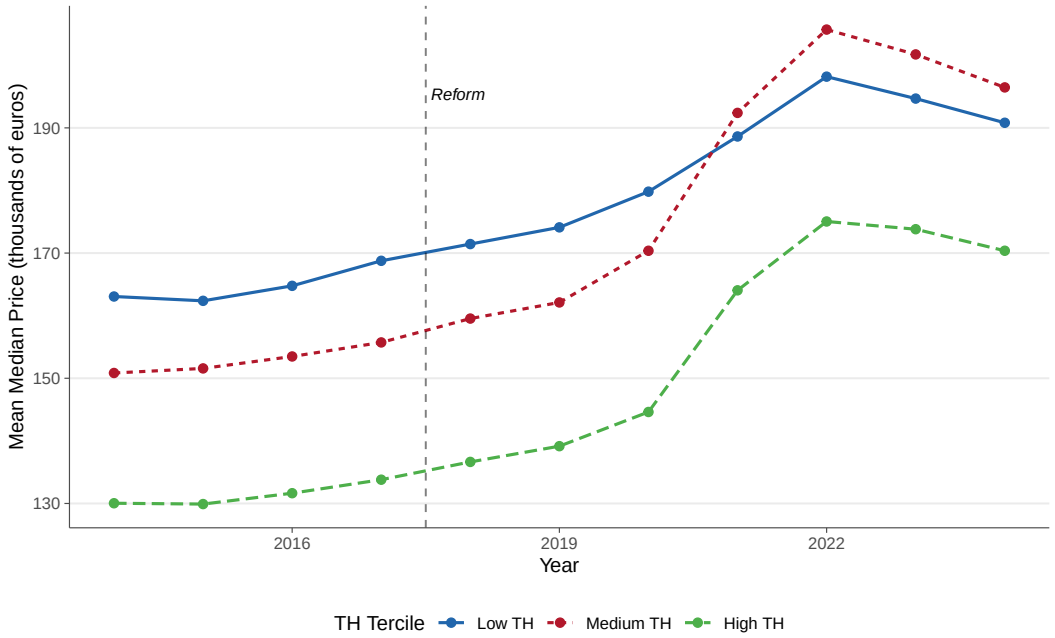

Figure 7: Raw Price Trends by TH Tercile, 2014–2024

Notes: Mean of commune-level median property prices (in thousands of euros) by TH tercile (based on 2017 TH rate). Dashed line indicates reform onset (2018). See **??** for the normalized version.

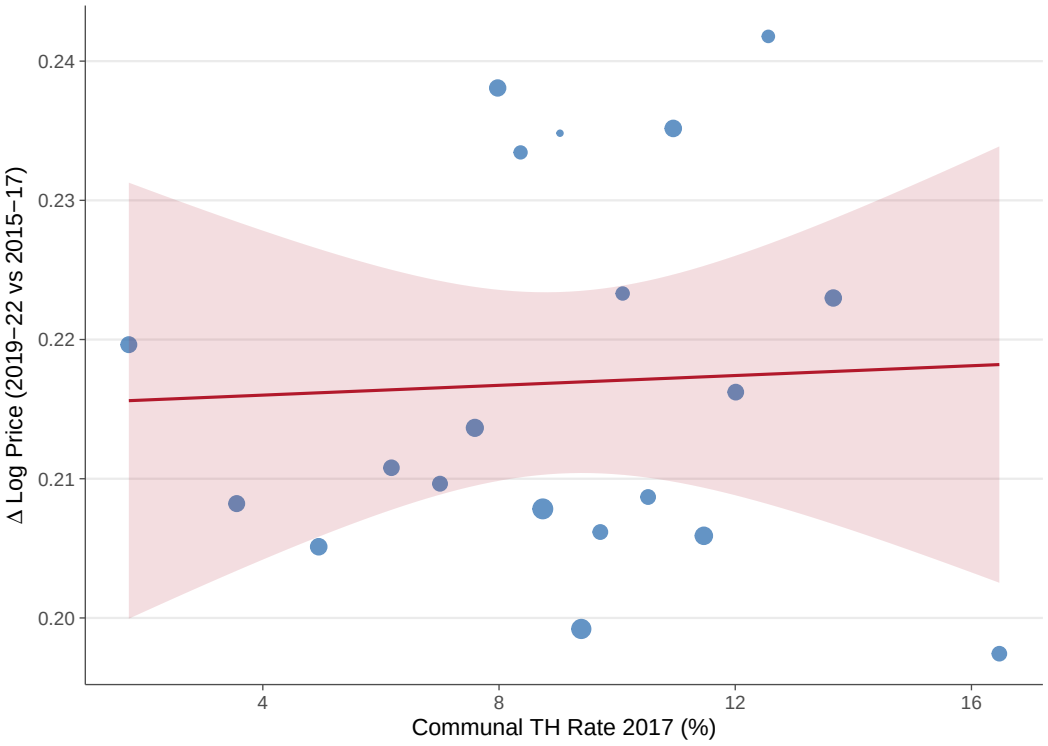

Figure 8: Binned Scatter: TH Rate and Price Change

Notes: Each point represents the mean of a 5-percentile bin of the 2017 communal TH rate. The y-axis shows the change in log median price between the pre-reform (2015–2017) and post-reform (2019–2022) periods. Point size proportional to number of communes. Red line is the OLS fit.