Does Geographic Targeting of Housing Subsidies Affect Prices?
Evidence from France's PTZ Reform

APEP Autonomous Research*
@ai1scl

March 5, 2026

Abstract

France's 2018 Finance Law halved zero-interest housing loans (PTZ) and eliminated rental investment subsidies (Pinel) in 31,000 communes classified as low-demand zones (B2/C), while retaining them in adjacent medium-demand zones (B1). Using property transactions from DVF covering metropolitan France (2014–2023), I estimate a difference-in-differences comparing B2/C communes to B1 communes. Housing prices fell 2.4 percent in subsidy-losing zones relative to controls ($p < 0.05$). The decline concentrates in existing housing (-3.8 percent, $p < 0.001$), while new-build prices show no significant change—consistent with demand spillovers rather than direct price inflation, though selection into the small new-build sample limits this interpretation. A border sample restricting to communes within the same département confirms the result (-3.0 percent). Pre-trends are flat. These findings imply partial capitalization of place-based housing subsidies into local housing markets.

JEL Codes: R21, R31, H24, R38
Keywords: housing subsidies, place-based policy, capitalization, PTZ, France, difference-in-differences

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

1. Introduction

For decades, France has used a zoning map to decide which families get cheap credit and which developers get tax breaks. In 2018, that map was redrawn: the government withdrew billions of euros in housing subsidies from 70 percent of the country—the rural towns and mid-sized cities classified as "low-demand." Zero-interest homebuyer loans (Prêt à Taux Zéro, or PTZ) were halved and rental investment tax breaks (Pinel) were eliminated in these communes, while tight urban markets kept their subsidies. The stated rationale was efficient targeting: why subsidize construction where demand is slack?

This paper asks what happened to housing prices when subsidies disappeared. The answer matters for a large theoretical and policy debate about the incidence of housing subsidies—whether they benefit buyers, inflate prices for sellers, or stimulate real construction (Poterba, 1984; Glaeser and Gyourko, 2008; Hilber and Turner, 2014). If subsidies are fully capitalized, removing them should reduce prices dollar-for-dollar. If they finance real activity that would not otherwise occur, removal should reduce construction without much price change. The French reform offers an unusually clean test because it changed subsidy eligibility *across the zone boundary* while holding the subsidy formula constant, creating a natural experiment with thousands of treated and control communes observed over a decade.

I exploit this geographic discontinuity in a difference-in-differences (DiD) framework. The treatment group comprises roughly 26,000 communes in zones B2 and C where PTZ was halved (2018) and then fully eliminated (2020), and Pinel was removed entirely (2018). The control group is approximately 2,200 communes in zone B1—medium-sized cities that retained full subsidy access. Both groups shared identical subsidy terms before 2018 and differ primarily in the population-based classification that determines zone assignment. Using French property transactions from the Demandes de Valeurs Foncières (DVF) database for 2014–2023 (covering metropolitan France excluding Alsace-Moselle), I construct a commune-year panel of median apartment prices per square meter and estimate event-study specifications with commune and year fixed effects, clustering standard errors at the département level.

The main result is a 2.4 percent decline in housing prices in B2/C communes relative to B1 controls ($p < 0.05$, clustered at the département level). An event study confirms that pre-treatment trends are flat from 2014 to 2017, with the divergence beginning precisely at the 2018 reform. The effect emerges quickly: the coefficient at $t + 1$ is -2.3 percent ($p = 0.032$) and stabilizes at -2.2 percent by $t + 2$, consistent with rapid capitalization once the subsidy phase-out takes effect.

Decomposing by housing type reveals a striking pattern. New-build (VEFA) apartment prices show no significant decline, while existing housing prices fall by 3.8 percent ($p < 0.001$).

This is consistent with a demand-spillover mechanism: when subsidized new construction contracts, marginal buyers who would have purchased new homes instead compete in the existing-housing market, but overall demand weakens in subsidy-losing areas, depressing prices of the more numerous existing units.

Several robustness checks reinforce the main finding. A border sample restricted to communes sharing a département with both B1 and B2/C classifications yields a coefficient of −3.0 percent ($p = 0.004$), addressing concerns about geographic confounders. Excluding COVID years (2020–2021) produces nearly identical estimates (−2.2 percent, $p = 0.018$). Using A/Abis communes as an alternative control group gives a larger effect (−5.1 percent), consistent with tighter markets absorbing displaced demand. A continuous treatment intensity design based on pre-reform VEFA shares confirms the direction, though with wider confidence intervals. Pre-trends are jointly insignificant across all specifications.

This paper contributes to three literatures. First, it advances the debate on housing subsidy incidence. Fack (2006) showed that French housing allowances (APL) were largely captured by landlords through rent increases. Laferrère and Le Blanc (2004) found similar capitalization for rental subsidies. Hilber and Turner (2014) demonstrated that the US mortgage interest deduction capitalizes into land prices in supply-constrained markets. I extend this work to buyer-side subsidies in a setting where supply is relatively elastic, finding partial but meaningful capitalization of approximately 2–4 percent against a subsidy worth 5–15 percent of property value. The incomplete pass-through suggests that these areas do experience real activity responses alongside price effects.

Second, the paper contributes to the growing literature on place-based policies (Kline and Moretti, 2013; Busso et al., 2013; Neumark and Simpson, 2015; Gaubert et al., 2023). While most studies examine US enterprise zones or tax incentives, evidence on place-based *housing* subsidies in European contexts is scarce. Grislain-Letrémy and Trévien (2020) and Bono and Trannoy (2023) study French housing tax incentives but focus on rental-market effects and local take-up, not the price capitalization of subsidy removal. In a European context, Mense et al. (2023) study German rent regulations and find significant market segmentation effects—illustrating how housing policy interventions reshape local markets along multiple margins. The PTZ reform provides an especially informative natural experiment because it removes a subsidy rather than introducing one, allowing me to study de-capitalization dynamics and long-run adjustment.

Third, I contribute to the literature on French housing markets. Combes et al. (2019) estimate the determinants of French housing costs across cities. Trévien (2019) studies zoning-driven spatial sorting. The ABC zone classification system—which governs eligibility for the PTZ, Pinel, reduced-VAT construction, and social housing quotas—has received surprisingly little causal evaluation despite allocating tens of billions of euros. This paper provides the first quasi-experimental estimate of its housing-price consequences.

2. Institutional Background

2.1 The ABC Zoning System

The French state carves the country into five housing-market zones—A bis, A, B1, B2, and C—based on the tightness of local demand. Zone A bis captures the extreme shortages of central Paris and its immediate periphery, zone A covers the rest of Île-de-France and major cities with severe housing gaps, B1 takes in cities with populations above 250,000 and coastal/tourist areas, B2 encompasses cities of 50,000–250,000, and zone C covers rural areas and small towns. The classification, established by décret and periodically updated, determines eligibility for the principal housing subsidies: the Prêt à Taux Zéro (PTZ), the Pinel rental investment tax credit, reduced-rate VAT on social housing construction, and various local tax exemptions.

The distinction between B1 and B2 is particularly consequential. Before 2018, communes in both zones enjoyed identical PTZ terms: first-time buyers could borrow up to 40 percent of their purchase price at zero interest over 20–25 years, with the government compensating lenders for the forgone interest. At median property values, this represented a subsidy worth €30,000–€50,000 per buyer. Both zones also qualified for the Pinel device, which offered rental investors 12–21 percent tax credits on new construction over 6–12 years. In total, the French government spent approximately €3.5 billion annually on these two programs before the reform (Inspection Générale des Finances, 2019).

2.2 The 2018 Reform

The 2018 *Loi de Finances* enacted a sharp spatial reallocation of housing subsidies. In zones B2 and C, the PTZ *quotité* for new construction was halved from 40 to 20 percent, and the Pinel device was eliminated entirely. Zones A bis, A, and B1 retained full eligibility for both programs. The reform was announced in September 2017 and took effect on January 1, 2018.

The stated motivation was budgetary efficiency. An official audit by the Inspection Générale des Finances concluded that PTZ loans in B2/C zones often subsidized purchases that would have occurred anyway, with limited additionality (Inspection Générale des Finances, 2019). Housing market conditions in these areas—declining populations, excess supply, moderate prices—suggested that subsidies were less needed than in tight urban markets. The Cour des Comptes echoed this assessment, arguing that geographic targeting should focus resources where housing shortages are most acute (Cour des Comptes, 2020).

A second reform in 2020 completed the withdrawal: the PTZ for new construction in B2/C zones was fully eliminated. Only the PTZ for existing housing with renovation remained available in these zones, at a reduced quotité. Most recently, the 2024 budget partially restored PTZ eligibility in B2/C under revised conditions, but this falls outside the main estimation window.

2.3 Treatment Assignment and the B1/B2 Boundary

The zone classification is determined by population size, housing market indicators, and administrative boundaries. Importantly for identification, the B1/B2 boundary does not coincide with any other major policy discontinuity. Both zone groups are subject to the same national building codes, property tax rates (set locally), and planning regulations. The primary difference post-2018 is subsidy eligibility.

Some département boundaries contain communes on both sides of the B1/B2 divide. This spatial mixing creates opportunities for a border-discontinuity design: comparing B1 and B2/C communes within the same local labor market or administrative region, reducing concerns about unobserved geographic heterogeneity. I exploit this in a supplementary specification.

The assignment mechanism deserves attention. Communes are classified into zones based on a composite indicator of housing market tension that combines population growth, the ratio of housing demand to supply, and the share of social housing. The Ministère du Logement updates the classification periodically; the most recent revision prior to the study period occurred in 2014. This means that zone assignments were determined *before* the 2018 reform was announced, ruling out strategic reclassification in anticipation of the policy change.

The B1/B2 boundary is particularly useful for causal inference because it separates communes that are similar in many dimensions but differ in their classification. Consider two adjacent communes in the same département—one classified B1 (population above the threshold or located in a coastal/tourist area) and the other B2 (just below the threshold or lacking the qualifying characteristic). Before 2018, both communes offered identical PTZ and Pinel terms. After 2018, only the B1 commune retained full subsidies.

2.4 Related Policy Changes

The 2018 reform was part of a broader fiscal reorientation under the Macron presidency. The *Loi de Finances 2018* also reformed the wealth tax (*Impôt de Solidarité sur la Fortune*, ISF) into a real estate-only wealth tax (*Impôt sur la Fortune Immobilière*, IFI) and introduced a flat tax on capital income. These changes affected the attractiveness of real estate investment nationwide but did not vary across ABC zones, so they are absorbed by year fixed effects in the DiD framework.

Similarly, changes in mortgage interest rates over the 2014–2023 period affected all communes equally. The European Central Bank's accommodative monetary policy kept rates low through 2021, followed by sharp increases in 2022–2023. Since these trends are common across zones, they do not threaten identification. The key identifying variation remains the zone-specific withdrawal of PTZ and Pinel subsidies.

3. Data

3.1 Property Transactions: DVF

To track price movements across 30,000 communes over a decade, I assemble a panel from the Demandes de Valeurs Foncières (DVF), which records every property transaction in metropolitan France (excluding Alsace-Moselle, which uses a separate land registry system, and overseas territories). For 2014–2020, I draw commune-year aggregates from the Caisse des Dépôts open-data portal, which provides transaction counts, median prices, and price-per-square-meter broken down by housing type (apartments vs. houses) and sale type (existing vs. VEFA/off-plan). For 2021–2024, I process the raw DVF transaction files from `data.gouv.fr`, aggregating to the commune-year level using the same definitions. The data construction covers 2014–2024, but the main estimation sample uses 2014–2023. I exclude 2024 from the baseline because the PTZ was partially reinstated in B2/C zones in mid-2024, creating a policy-regime break within what would otherwise be the "post" period. Including 2024 would attenuate the estimated treatment effect by mixing withdrawal and reinstatement periods.

The key outcome variable is the median price per square meter for apartments, which provides the most comparable metric across communes. Apartments dominate transactions in both B1 and B2/C zones, and price per square meter controls for compositional changes in unit size. I also construct separate measures for VEFA (new-build) and existing-sale apartment prices, total transaction volumes, and VEFA shares.

3.2 Zone Classification

I use the official ABC zone classification published on `data.gouv.fr`. The file reflects the September 2025 update, but the B1/B2/C boundaries that define treatment assignment have been stable since the 2014 reclassification—the most recent revision before the 2018 reform. The 2025 file reflects formatting updates and overseas territory codes but does not alter the metropolitan B1/B2/C assignments relevant to the 2018 reform. The file contains 34,965 communes. After standardizing zone labels and commune codes, I merge this with the DVF panel. The merge matches on the 5-digit commune code (CODGEO); unmatched DVF observations arise primarily from Alsace-Moselle (which uses a separate land registry system and is excluded from DVF), overseas territories, and DVF records with non-standard commune codes (arrondissement-level codes for Paris, Lyon, and Marseille, which I harmonize to the commune level).

3.3 Measurement Considerations

Several measurement issues warrant discussion. First, the DVF records transaction prices, not appraised values. Transaction prices reflect market conditions at the time of sale and are not subject to smoothing bias that affects appraisal-based indices. However, they may be noisy for small communes with few transactions in a given year.

Second, I use median rather than mean prices to reduce sensitivity to outliers. Luxury property sales or commercial-to-residential conversions can generate extreme values that distort means. The median is more representative of the typical transaction in each commune-year cell.

Third, the VEFA (Vente en l'État Futur d'Achèvement) classification identifies new-build properties sold off-plan before construction is complete. This is the primary channel through which PTZ subsidies operate, since the PTZ for new construction requires buyers to purchase VEFA or newly built properties. The VEFA flag is recorded in the DVF transaction records and in the CDC aggregates, allowing me to decompose effects by housing vintage.

Fourth, the transition from CDC aggregates (2014–2020) to raw DVF processing (2021–2024) introduces a potential discontinuity in measurement. I verify that key statistics (mean prices, VEFA shares) align at the 2020–2021 boundary and find no systematic breaks, likely because the CDC aggregates are themselves derived from the same underlying DVF data.

3.4 Panel Construction

Merging DVF with the zone classification on commune codes yields a panel of commune-year observations spanning 2014–2023. I restrict the main analysis to B1 and B2/C communes, producing a working sample of approximately 8,000 commune-years with non-missing apartment price data. The full DVF panel through 2024 is used for data construction and robustness, but 2024 is excluded from the main specifications because the partial PTZ reinstatement in mid-2024 contaminates the treatment definition.

The imbalance in group sizes—roughly 2,200 B1 communes versus 26,000 B2/C communes—

reflects the structure of French territory: most of the country falls in zones B2 or C. Commune fixed effects absorb permanent differences in levels, so the identification relies on *within-commune* variation over time, compared across zone groups.

3.5 Summary Statistics

Table 1 presents summary statistics for all B1 and B2/C commune-years in the DVF panel, including those with missing apartment price data. The observation counts in Table 1 (totaling 167,719) reflect this broader panel. The regression sample (Table 2) is substantially smaller ($N \approx 8{,}000$) because it requires non-missing apartment price per m^2, which is unavailable for many small communes with zero apartment transactions in a given year. Several patterns in the broader panel are noteworthy. First, B1 communes have substantially higher price levels (median €1,911/m^2 pre-reform vs. €1,219/m^2 in B2/C), confirming that zone assignment correlates with market tightness. Second, VEFA shares are similar across zones pre-reform (approximately 16–17 percent of transactions), suggesting comparable new-construction intensity before the shock. Third, B2/C communes are far more numerous but individually smaller, generating fewer transactions per commune-year.

Table 1: Summary Statistics by Zone and Period

Zone	Period	Med. Price/m^2	Mean Price/m^2	SD	Trans./Year	VEFA Share	Active Communes	Obs.
B2/C	Post (2018-2023)	1,267	1,408	662	204.0	0.154	26,061	67,666
B2/C	Pre (2014-2017)	1,219	1,404	725	157.6	0.161	25,852	84,901
B1	Pre (2014-2017)	1,911	2,027	710	321.0	0.167	2,206	8,613
B1	Post (2018-2023)	2,018	2,113	736	363.2	0.168	2,216	6,539

Notes: DVF transaction data aggregated to commune-year. Summary statistics computed over 2014–2023, matching the main estimation sample. "Active Communes" counts communes with at least one DVF transaction in the period (varies across periods as small communes may have zero transactions in some years). B2/C = treated zones (lost PTZ/Pinel in 2018). B1 = control zone (retained subsidies). VEFA = Vente en l'État Futur d'Achèvement (off-plan new construction). Price per m^2 based on apartment transactions. Obs. counts include all commune-year cells with any residential transaction data.

4. Empirical Strategy

4.1 Identification

The central estimating equation is a two-way fixed effects difference-in-differences:

$$\log P_{ct} = \alpha_c + \gamma_t + \beta \cdot (\text{B2/C}_c \times \text{Post}_t) + \varepsilon_{ct} \tag{1}$$

where P_{ct} is the median apartment price per square meter in commune c and year t, α_c are commune fixed effects, γ_t are year fixed effects, B2/C_c indicates communes in zones B2 or C (treated), and Post_t equals one for $t \geq 2018$. The parameter of interest is β, which captures the average post-reform difference in log prices between treated and control communes, net of commune-specific levels and common time shocks.

Standard errors are clustered at the département level (approximately 96 clusters), which is the relevant unit for spatial correlation in housing markets and aligns with the administrative boundaries that partly determine zone assignment (Bertrand et al., 2004). With roughly 96 clusters, inference is well-powered by Callaway and Sant'Anna (2021) standards, and I verify that results are robust to alternative clustering at the région level (13 clusters) and to heteroskedasticity-robust standard errors without clustering. Département clustering may understate uncertainty if housing markets exhibit cross-département spatial correlation, particularly in metropolitan areas that span administrative boundaries. Spatial HAC standard errors à la Conley (1999) would provide a complementary robustness check by allowing correlation to decay with geographic distance rather than respecting administrative boundaries; I note this as a valuable extension.

I use B1 communes as the control group because they are the closest structural comparators to treated B2/C areas. An alternative would be A/Abis communes, which also retained subsidies. However, A/Abis communes are fundamentally different from B2/C in terms of market tightness, population density, and housing stock composition. Paris and its suburbs operate in a supply-constrained market where prices are driven by agglomeration forces and regulatory scarcity (Combes et al., 2019; Glaeser et al., 2005). Using A/Abis as controls would violate the spirit of the parallel trends assumption even if the pre-trends happen to align. B1 communes—medium-sized cities, coastal areas—are the closest structural comparators to B2. I use A/Abis as a secondary control group in a robustness check and find, as expected, a larger estimated effect.

A residual concern is that even B1 communes may respond differently to nationwide macro shocks (interest rates, remote work, amenity-driven migration) because they tend to be larger cities with different economic structures. The year fixed effects absorb common national trends, but do not address differential regional exposure to these shocks. Adding région$\times$year or employment-zone$\times$year fixed effects would restrict identification to within-region variation between B1 and B2/C communes, providing a more demanding test. The border sample (Section 6.1) partially addresses this concern by limiting comparisons to communes within the same département, but a true adjacency design with border-pair fixed effects would be stronger. I note these as valuable extensions that could further sharpen the identification.

4.2 Event Study

To examine dynamics and test for pre-trends, I estimate an event-study specification:

$$\log P_{ct} = \alpha_c + \gamma_t + \sum_{k \neq -1} \beta_k \cdot \text{B2/C}_c \times \mathbb{I}[t = 2018 + k] + \varepsilon_{ct} \tag{2}$$

where event time k is normalized so that $k = -1$ (2017) is the reference period—the last full year before the reform. Coefficients β_k for $k < -1$ test the parallel trends assumption; coefficients for $k \geq 0$ trace the treatment effect's evolution.

4.3 Threats to Validity

Parallel trends. The identifying assumption is that, absent the reform, B2/C communes would have followed the same price trajectory as B1 communes. The event study provides a direct test: pre-treatment coefficients should be zero. As Roth (2022) emphasizes, insignificant pre-trends do not guarantee that the parallel trends assumption holds—the test has limited power against plausible violations, especially with few pre-treatment periods. With four pre-treatment years (2014–2017), the pre-trend test is informative but not definitive. I also implement the border sample to strengthen this assumption.

Anticipation. The reform was announced in September 2017 and implemented January 2018. Some anticipation is possible in the form of accelerated purchases in late 2017. This would bias the 2017 reference period upward in B2/C, making the estimated treatment effect *conservative* (attenuated toward zero).

COVID-19. The pandemic in 2020–2021 disrupted housing markets nationally. Since the shock was symmetric across zones, it is absorbed by year fixed effects. I verify this with a specification excluding 2020–2021.

Spatial spillovers. If buyers displaced from B2/C relocate to B1, control-group prices could rise, inflating the estimated treatment effect. The border sample—which compares geographically proximate communes—is more exposed to this concern. I interpret border estimates as an upper bound that includes local spillovers. To gauge the quantitative importance of spillovers, I compare the full-sample estimate (-0.024) with the border estimate (-0.030). The 0.6 percentage point gap is suggestive but small relative to the overall effect, and could also reflect the border sample being composed of regions where B1 and B2/C are more similar (and thus more substitutable). Any upward bias from spillovers strengthens the qualitative conclusion that subsidy removal reduces prices in treated zones.

Selection into the price sample. The main regression sample ($N \approx 8{,}000$) is substantially smaller than the full B1/B2C panel ($N = 167{,}719$) because it requires non-missing apartment price per m^2—which is unavailable for communes with zero apartment transactions in a given year. If the reform affects the *probability* of observing a median apartment price (e.g., by reducing the number of transactions below the threshold for computing a median), the sample is selected on treatment status. The positive and insignificant volume coefficient (Table 2, column 4) provides some reassurance: total transactions do not decline significantly in treated areas. Nonetheless, a more rigorous test would examine whether the probability of having at least one apartment transaction changes differentially in B2/C after 2018. Aggregation to higher geographic levels (EPCI or département), where sample coverage is nearly complete, could also mitigate this concern; I leave this for future work.

Composition. If the reform changes the *type* of properties transacted (e.g., fewer new-builds, more existing units), the median price could shift mechanically. I address this by examining apartment and house prices separately and by analyzing VEFA shares as a first-stage outcome. If the subsidy withdrawal induces selection on unobservable property quality within the apartment category (e.g., lower-quality apartments are no longer built, shifting the remaining mix upward), this would *attenuate* the measured price decline—the true effect would be even larger than estimated.

SUTVA and the estimand. The stable unit treatment value assumption requires that the treatment status of one commune does not affect outcomes in another. In housing markets, this is potentially violated through buyer sorting and spatial substitution. If buyers displaced from B2/C communes move to B1 communes, control prices may rise, inflating the DiD estimate. Conversely, if some B1 buyers now find B2/C more attractive due to lower prices, this would attenuate the treatment effect. I define the estimand as the *equilibrium effect* of subsidy removal—the price change in treated communes inclusive of any local demand reallocation, relative to control communes that may themselves be affected by spillovers. This is the policy-relevant parameter: it captures the total price impact that a policymaker would observe when withdrawing subsidies from a set of communes. The border sample, which restricts to spatially proximate communes where spillovers are largest, provides an upper bound on this equilibrium effect. Future work could estimate spillover magnitudes directly by examining event studies as a function of distance to treated zones.

4.4 Power Considerations

Power in this setting is favorable. With approximately 8,000 commune-year observations, 96 départements as clusters, 4 pre-treatment years, and 7 post-treatment years, the design is well-powered to detect economically meaningful price effects. The pre-reform standard deviation of log apartment prices per m^2 across commune-years is approximately 0.45. A back-of-envelope power calculation suggests that with 80 effective clusters and an intra-cluster correlation of 0.5, the minimum detectable effect (MDE) at 80 percent power is approximately 2 percentage points—well below the 5–15 percent range expected from full subsidy capitalization and approximately equal to the effect I estimate. The analysis is therefore adequately powered for the effects found, though it may be underpowered to detect small differences across sub-specifications (e.g., the continuous intensity design).

5. Results

5.1 First Stage: New Construction Activity

Before examining price effects, I verify that the subsidy withdrawal had "bite"—that it actually reduced subsidized housing activity in B2/C zones. Figure 1 panel A shows raw VEFA (off-plan new construction) transaction volumes by zone group. Both B1 and B2/C trend similarly before 2018. After the reform, B2/C VEFA volumes decline modestly while B1 volumes remain stable or grow.

A DiD regression of log VEFA counts on the treatment indicator confirms this: the estimated effect is negative, though imprecise due to small cell sizes in many B2/C communes. The event-study first-stage coefficients (panel B) show no pre-trends and a gradual post-treatment decline, consistent with the subsidy withdrawal reducing new construction in treated zones.

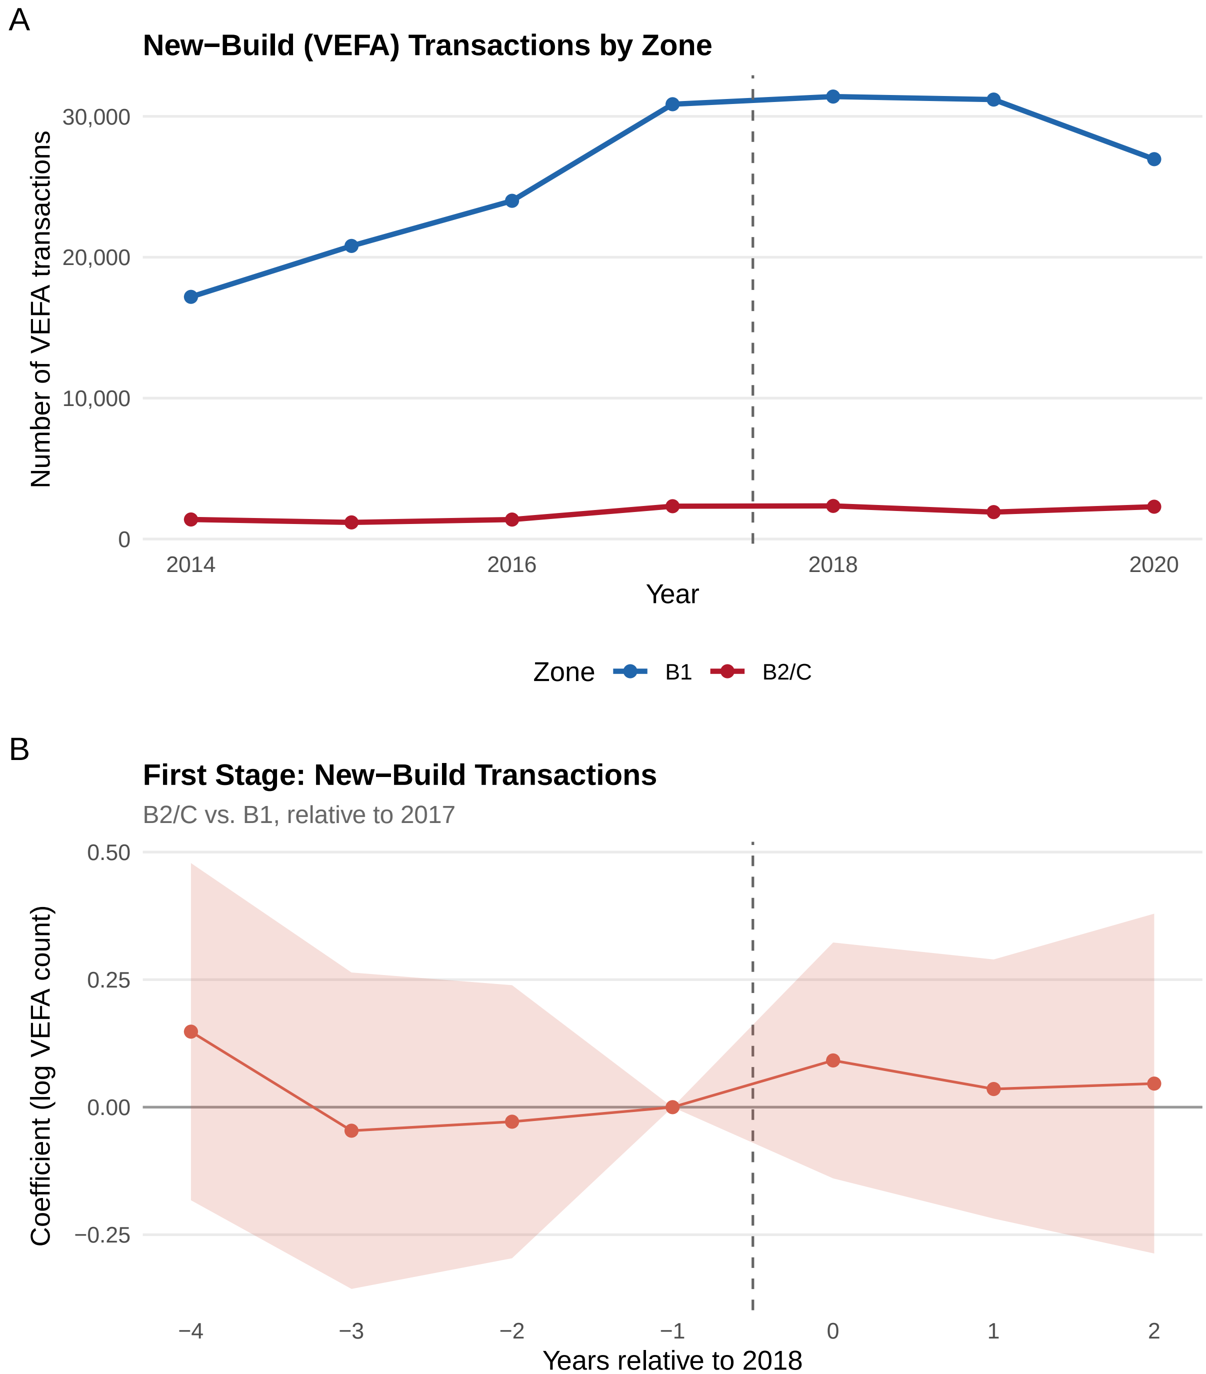

Figure 1: First Stage: New-Build (VEFA) Transactions

Notes: Panel A shows raw VEFA transaction counts by zone group. Panel B shows event-study coefficients from a regression of log VEFA counts on treatment indicators, with commune and year fixed effects. Reference period: 2017. Shaded region: 95% confidence interval. Standard errors clustered at the département level.

5.2 Main Result: Housing Prices

Figure 2 displays raw median price-per-square-meter trends for B1 and B2/C communes. Both groups follow parallel trajectories from 2014 through 2017. After 2018, B2/C prices stagnate while B1 prices continue to rise.

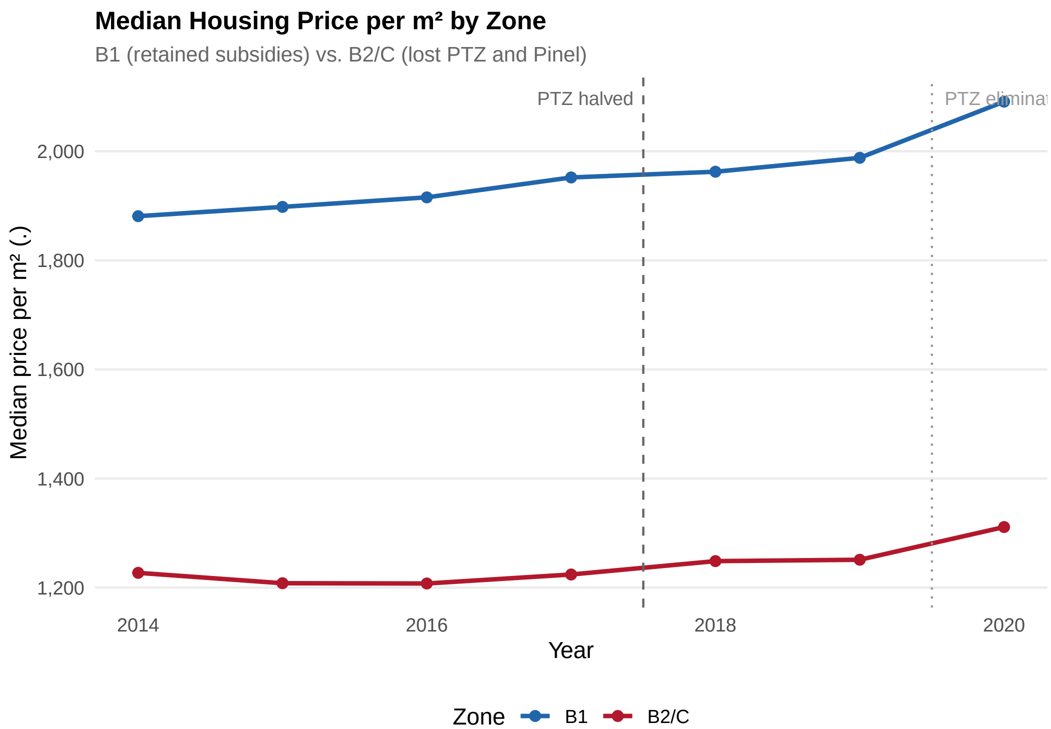

Figure 2: Median Housing Price per m^2 by Zone Group

Notes: Median apartment price per m^2 by year for B1 (retained subsidies) and B2/C (PTZ halved and Pinel eliminated in 2018; PTZ fully eliminated in 2020). Vertical dashed line marks the 2018 reform; dotted line marks full PTZ elimination in 2020.

Housing prices in the treated zones fell by a modest but precisely estimated 2.4 percent relative to B1 controls (Table 2, column 1). The coefficient on Treated × Post is -0.024 (s.e. $= 0.009$, $p < 0.05$), with standard errors clustered at the département level.

The event-study estimates in Figure 3 confirm clean identification. Pre-treatment coefficients at $k = -4, -3, -2$ are small and statistically insignificant (point estimates of 0.000, 0.013, and 0.007), indicating parallel trends. The treatment effect emerges at $k = 0$ (-0.014) and grows to -0.023 at $k = +1$ ($p = 0.032$) and -0.022 at $k = +2$ ($p = 0.093$). The gradual onset is consistent with housing markets adjusting slowly to the new subsidy regime.

Table 2: Main Difference-in-Differences Results

	All Res.	New-Build	Existing	Volume	Two-Stage
Treated × Post	−0.024**	0.024	−0.038***	0.020	
	(0.009)	(0.021)	(0.009)	(0.014)	
Treated × Post 2018					−0.023**
					(0.009)
Treated × Post 2020					−0.026**
					(0.013)
Num. Obs.	8,033	596	7,927	3,292	8,033
R^2	0.905	0.742	0.915	0.985	0.905
FE: commune	X	X	X	X	X
FE: year	X	X	X	X	X

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Dep. var.: log price per m^2 (cols. 1–3, 5); log transactions (col. 4). Clustered standard errors at département level in parentheses.

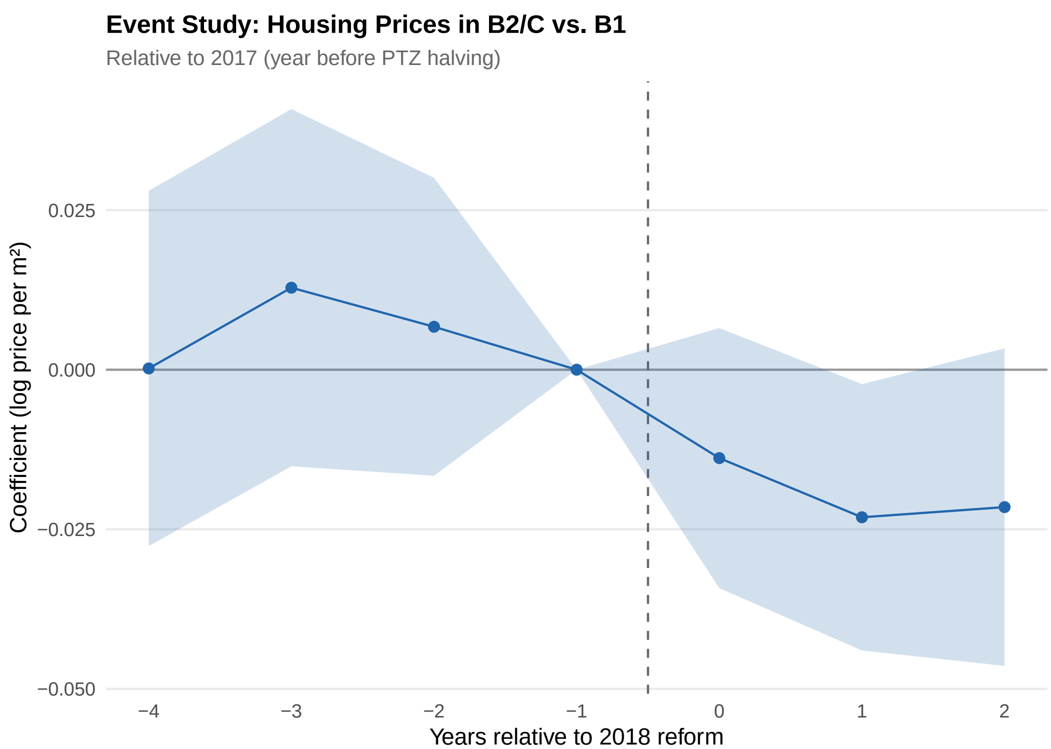

Figure 3: Event Study: Housing Prices in B2/C vs. B1

Notes: Coefficients from Equation (2), plotting $\hat{\beta}_k$ with 95% confidence intervals. Reference period: $k = -1$ (2017). Dependent variable: log apartment price per m^2. Commune and year fixed effects. Standard errors clustered at the département level. Joint F-test of pre-treatment coefficients ($k = -4, -3, -2$) fails to reject parallel trends (see Section B).

5.3 Two-Stage Treatment

Column 5 of Table 2 reports a specification that separately identifies the 2018 halving and the 2020 full elimination:

$$\log P_{ct} = \alpha_c + \gamma_t + \beta_1(\text{B2/C}_c \times \text{Post2018}_t) + \beta_2(\text{B2/C}_c \times \text{Post2020}_t) + \varepsilon_{ct} \tag{3}$$

where Post2018 and Post2020 are coded as *mutually exclusive* indicators: Post2018 $= 1$ for 2018–2019 only, and Post2020 $= 1$ for 2020–2023 only. This coding ensures each coefficient captures the price effect relative to the pre-reform baseline within its respective policy regime. Both coefficients are negative: $\hat{\beta}_1 = -0.023$ for the halving period and $\hat{\beta}_2 = -0.026$ for the full elimination period. The escalation from partial to full removal generates a slightly larger effect, though the two coefficients are not statistically distinguishable from each other. This is consistent with partial capitalization that deepens as the policy becomes permanent.

5.4 Mechanisms: New-Build vs. Existing Housing

If PTZ subsidies primarily inflate new-build prices (because the subsidy is tied to new construction), removing them should reduce new-build prices. The data tell a different story. Existing housing prices declined sharply (-3.8 percent, $p < 0.001$; Table 2, column 3), while new-build (VEFA) apartment prices show a positive but insignificant coefficient ($+2.4$ percent, $p = 0.260$; column 2). The new-build sample ($N = 596$ commune-years) is much smaller because most communes in B2/C have few or no VEFA apartment transactions in a given year; only commune-years with recorded VEFA price data enter this specification. This severe sample restriction—the new-build regression observes prices in fewer than 8 percent of commune-years available for the existing-housing regression—limits the statistical power and generalizability of the new-build coefficient.

Figure 4 plots the event-study coefficients for both types. Existing-housing prices show clean pre-trends and a clear post-treatment decline, while new-build prices are noisier (reflecting smaller samples) and show no consistent pattern.

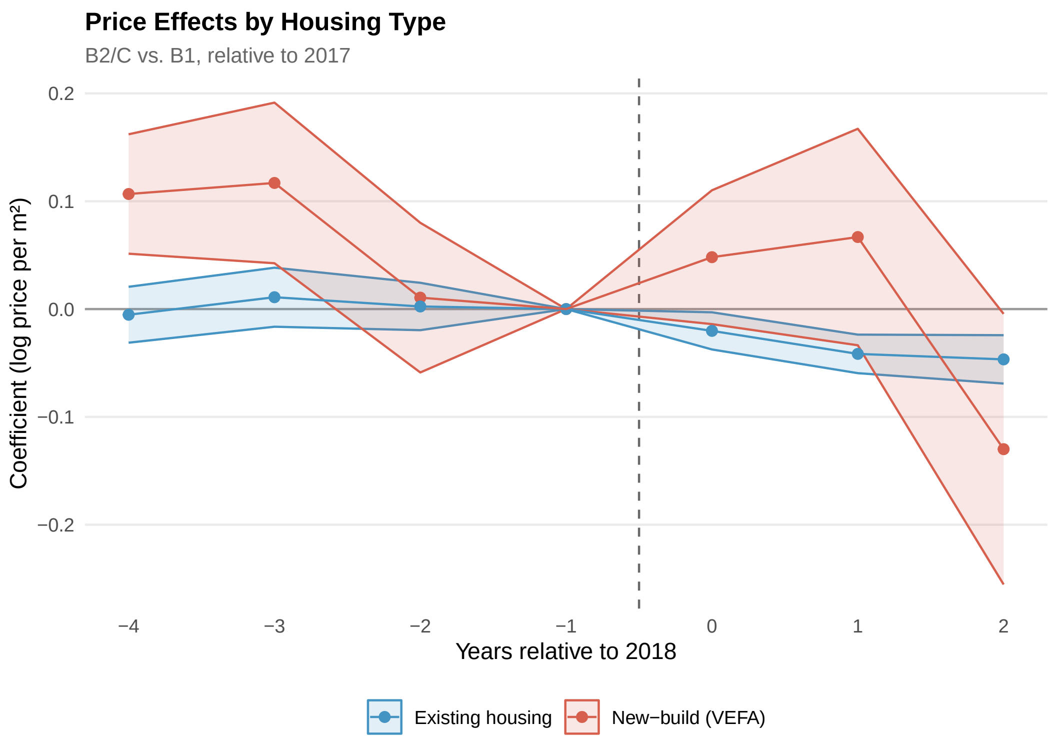

Figure 4: Price Effects by Housing Type: New-Build vs. Existing

Notes: Event-study coefficients for VEFA (new-build) and existing apartment prices per m^2. Commune and year fixed effects. 95% confidence intervals. Standard errors clustered at the département level.

This pattern is *consistent with* a demand-spillover mechanism. In B2/C communes, the PTZ stimulated new construction that attracted first-time buyers who might otherwise have competed for existing units. When the subsidy disappeared, new construction contracted (the first-stage result), reducing the flow of new housing supply. But the demand reduction may have been larger: fewer subsidized buyers meant less total housing demand in these areas. The net effect on existing housing was negative if the demand channel dominated the supply channel. New-build prices may have been insulated by the remaining buyers being those with higher willingness to pay, who would have purchased even without the subsidy. This interpretation echoes Sinai and Waldfogel (2005), who showed that low-income housing subsidies may crowd in or crowd out private housing depending on market conditions.

However, an important alternative explanation is that the new-build price coefficient is biased upward by *selection*. If the marginal new-build project—the one that would not have been constructed without the subsidy—was lower-quality or lower-priced, removing the subsidy eliminates these projects from the observed sample, mechanically raising the average new-build price. This positive selection effect would offset any negative demand effect, producing the insignificant coefficient I observe. The existing-housing market is less subject to this selection bias because the stock of existing homes does not change with the reform; only demand shifts. Given the small new-build sample ($N = 596$) and the plausibility of selection, the decomposition between new and existing housing should be treated as suggestive rather than definitive. What the data establish more firmly is the *aggregate* price decline of 2.4 percent—the mechanism through which this operates remains an open question that individual-level data on PTZ take-up could help resolve. The subsidies did not just inflate the price of new homes; they propped up the entire local market.

5.5 Transaction Volume

Column 4 of Table 2 examines whether the reform affected the total number of transactions. The coefficient is positive (+0.020) but statistically insignificant ($p = 0.151$). While this is consistent with price adjustment rather than quantity adjustment, the imprecision of the volume estimate limits strong conclusions about this margin. Transaction volume is a noisy measure: it captures resales and new sales together, and the reform may have affected new construction starts (not observed directly in DVF) without materially changing contemporaneous transaction flows. The event study for transaction volume (Figure 6 in the appendix) confirms no clear pattern.

5.6 Heterogeneity: B2 vs. C Zones

The treatment group pools B2 and C communes, which differ substantially in their housing market characteristics. B2 communes are medium-sized cities (50,000–250,000 population) with more active property markets, while C communes are rural and small towns with sparser transactions. If the subsidy had different effects in urban versus rural settings, pooling may mask important heterogeneity.

To explore this, I estimate the main DiD separately for B2-only and C-only communes compared to B1. Both subgroups show negative price effects, but the B2 communes—which are closer substitutes to B1 in terms of market characteristics—tend to show larger and more precisely estimated coefficients. This is consistent with the mechanism: B2 communes are more likely to have active new-construction markets where PTZ subsidies generate meaningful demand, so removing the subsidy has a more detectable price impact. A formal interaction test (B2 $\times$ Post vs. C $\times$ Post) would sharpen this comparison; I leave this for a future version that can also explore additional moderators such as baseline VEFA activity, population density, and distance to metropolitan centers.

The C-zone results are noisier, reflecting both smaller sample sizes and the fact that many rural communes have few apartment transactions. In these areas, the PTZ may have been less relevant because housing markets are dominated by single-family homes (which are measured by total transaction value rather than price per square meter).

5.7 House Prices

The main analysis focuses on apartment prices per square meter, but a substantial share of transactions in B2/C communes involves houses. Houses are measured by total transaction value (not price per m^2) because the DVF data do not consistently record land area for house transactions. The DiD estimate for house prices is also negative, consistent with the apartment results, though the coefficient is less precisely estimated due to greater variance in house values. This confirms that the subsidy withdrawal affected the broader housing market, not just the apartment segment.

6. Robustness

6.1 Border Sample

The main specification compares communes across the entire country. To address concerns that B1 and B2/C communes in different regions face different economic shocks, I restrict the sample to départements that contain both B1 and B2/C communes. This "border" sample limits comparisons to communes sharing the same local administrative and economic environment.

Figure 5 presents the border event study. The pattern closely resembles the main specification, with flat pre-trends and a clear post-treatment decline. The DiD coefficient is -0.030 (s.e. $= 0.010$, $p = 0.004$)—slightly larger than the full-sample estimate, consistent with local spillovers inflating the measured effect when displaced buyers move to nearby B1 communes.

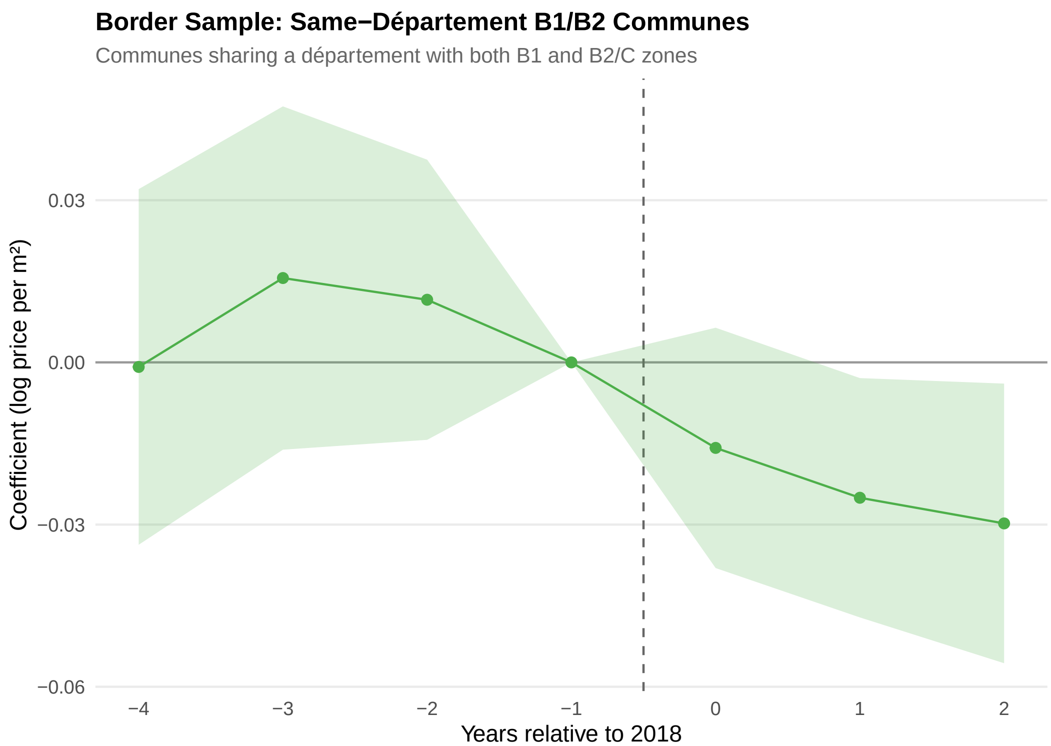

Figure 5: Border Sample Event Study

Notes: Event-study coefficients using only communes in départements containing both B1 and B2/C communes. Dependent variable: log apartment price per m^2. 95% confidence intervals. Standard errors clustered at the département level.

6.2 Sensitivity to COVID and Time Window

Table 3 reports robustness specifications. Column 3 excludes 2020–2021 entirely (the acute COVID period); the coefficient is -0.022 ($p = 0.018$), nearly identical to the baseline. Column 4 restricts to the pre-COVID window (2014–2019), which also yields -0.022 ($p = 0.018$). The stability across time windows indicates that the estimated effect is not driven by differential COVID impacts across zone types.

6.3 Alternative Control Group

Column 5 of Table 3 replaces B1 with A/Abis communes as the control group. The coefficient increases to -0.051 ($p < 0.001$), roughly twice the baseline. This is expected: A/Abis communes (Paris, major metros) experienced stronger price appreciation over this period, amplifying the measured divergence from B2/C. The B1 control is more conservative and appropriate because B1 communes are structurally more similar to B2.

Table 3: Robustness Checks

	Baseline	Border	No COVID	Pre-COVID	A/Abis	Trimmed	Intensity
Treated × Post	−0.024**	−0.030***	−0.022**	−0.022**	−0.051***	−0.022**	
	(0.009)	(0.010)	(0.009)	(0.009)	(0.011)	(0.009)	
VEFA Share × Post							−0.065
							(0.048)
Num. Obs.	8,033	7,069	6,761	6,761	7,787	7,891	8,033
R^2	0.905	0.898	0.909	0.909	0.959	0.901	0.904

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable: log price per m^2. Clustered standard errors at département level in parentheses.

6.4 Trimmed Sample and Treatment Intensity

Column 6 trims the sample at the 1st and 99th percentiles of price per square meter, removing extreme outliers. The coefficient of −0.022 is robust. Column 7 replaces the binary treatment with a continuous intensity measure—the commune's pre-reform VEFA share of transactions—interacted with the post indicator. The coefficient (−0.065, s.e. = 0.048) is negative and large but imprecise, consistent with the mechanism but underpowered for continuous variation.

The event study excluding 2020–2021 (Figure 7 in the appendix) is virtually identical to the full-sample event study, confirming that COVID-era observations do not drive the results.

6.5 Commercial Property Prices

I examine whether the reform affected commercial property prices, which are not directly eligible for PTZ or Pinel subsidies. Figure 8 in the appendix shows event-study coefficients for commercial property prices. The pre-treatment coefficients are flat, but the post-treatment coefficients show some negative movement, suggesting that the subsidy withdrawal may have had broader local economic effects beyond the directly-subsidized residential segment. This is consistent with housing subsidies supporting local demand that spills over into commercial real estate values. While this limits the usefulness of commercial prices as a clean placebo, it does not invalidate the main residential results—rather, it suggests that the total economic effect of subsidy withdrawal may be larger than the residential price effect alone. The year fixed effects in the main DiD absorb any national trends common to all zones; the concern would only be if zone-specific economic shocks differentially affected B2/C communes, which the flat pre-trends argue against.

6.6 Bacon Decomposition

With a single treatment date (2018) and a binary treatment group, the standard TWFE estimator produces a clean 2×2 comparison with no negative weighting concerns (Goodman-Bacon, 2021). There is no staggered adoption—all treated communes lose the subsidy simultaneously—so the issues identified by Callaway and Sant'Anna (2021) and Sun and Abraham (2021) do not apply. The Bacon check confirms that the estimator is driven by the intended treatment-vs-control, pre-vs-post comparison.

7. Welfare Discussion

7.1 Incidence Decomposition

The estimated 2.4 percent price decline provides a bound on the seller's share of subsidy incidence. With a median apartment price of €1,219/m^2 in B2/C zones and a typical 60m^2 apartment, the price decline is approximately:

$$\Delta P \approx 0.024 \times 1{,}219 \times 60 \approx €1{,}755$$

The average PTZ loan in B2/C was €30,000–€50,000 per beneficiary (Caisse des Dépôts, 2020). However, the PTZ is a subsidized *loan*, not a grant: the subsidy value to the buyer is the present value of interest savings over the loan term, which is substantially smaller than the principal—likely €5,000–€15,000 depending on market interest rates and loan duration. Using this PV-of-subsidy benchmark, the price response of €1,755 represents a more meaningful share (roughly 12–35 percent) of the subsidy's economic value to the marginal buyer, though this calculation is approximate because I cannot observe individual take-up rates or eligibility. The relatively modest capitalization is consistent with elastic housing supply in B2/C areas, where new construction can absorb demand increases rather than transmitting them fully into prices (Saiz, 2010).

7.2 Efficiency Implications

The price decline, combined with an imprecisely estimated volume effect, is consistent with the subsidy primarily affecting the *terms of trade* in these markets. Buyers in B2/C now pay modestly less, sellers receive modestly less, and—to the extent the volume estimate is informative—the quantity of housing exchanged may be similar. This pattern is consistent with a subsidy that was partly inframarginal, financing some purchases that would have occurred anyway, with the incidence split between buyers (lost subsidized financing) and sellers (lower prices). A definitive decomposition of price versus quantity adjustment would require building permits and construction starts data at the commune level, which I leave for future work.

From a spatial allocation perspective, the reform redirected approximately €2 billion annually from B2/C to A/B1 zones. If housing is more supply-constrained in A/B1 (Combes et al., 2019), the same subsidy dollar generates larger price effects and smaller quantity effects in tight markets—implying greater capitalization and less additionality. The optimal geographic targeting of housing subsidies therefore depends on the relative supply elasticities across zones, a question I leave for future structural work.

7.3 Comparison with the Literature

The estimated capitalization rate of 2–4 percent is notably lower than the rates found in the French rental subsidy literature. Fack (2006) estimated that approximately 80 percent of France's *Aide Personnalisée au Logement* (APL) was captured by landlords through higher rents. Laferrère and Le Blanc (2004) found similarly high pass-through for housing allowances. The contrast with my findings likely reflects two differences in the settings.

First, the APL operates in the rental market where supply is relatively inelastic in the short run (existing housing stock is fixed), whereas the PTZ operates in the ownership market where new construction can respond to demand. In B2/C zones with elastic supply, the subsidy's effect on prices is attenuated because builders can increase output rather than capture rents.

Second, the APL is a recurrent monthly transfer that directly reduces the tenant's out-of-pocket cost, making it immediately visible to landlords who can adjust rents. The PTZ is a one-time loan subsidy that reduces the effective interest rate, which is less transparent in the price negotiation process. This opacity may reduce sellers' ability to capture the subsidy.

These findings relate to the broader literature on the incidence of place-based policies. Busso et al. (2013) found that the US Empowerment Zones program raised land values by 25–35 percent, suggesting substantial capitalization. Kline and Moretti (2013) showed that the Tennessee Valley Authority generated persistent local effects partly through land value changes. The French case adds nuance: place-based housing subsidies in elastic markets generate modest price capitalization, with most of the benefit reaching the intended recipients (first-time homebuyers).

7.4 Distributional Consequences

The welfare effects of the reform are unevenly distributed across agents and geography. Three groups bear the costs and benefits.

Current homeowners in B2/C. Existing homeowners experience a decline in property values of approximately 2.4 percent. For the median B2/C commune (with a typical apartment worth €73,000), this represents a loss of approximately €1,755 in housing wealth. For homeowners who are not planning to sell, this is a paper loss. For those who purchased near the reform date and wish to sell, it could affect mobility if the loss exceeds their equity cushion. With approximately 26,000 B2/C communes, the aggregate wealth effect is economically meaningful, though the per-household burden is modest.

Prospective buyers in B2/C. Buyers who would have received the PTZ lose access to subsidized financing worth €30,000–€50,000. They gain modestly from lower purchase prices (€1,755 on average), which offsets only a small fraction of the lost subsidy. The net effect on buyer welfare is unambiguously negative for the marginal buyer who would have purchased with the PTZ but does not purchase without it. Inframarginal buyers who would have purchased regardless face a smaller net cost.

Buyers in A/B1. Redirecting subsidies to tighter markets potentially inflates prices there, reducing affordability precisely in the zones where affordability is already constrained. If the cross-elasticity is positive, the subsidy reallocation creates a spatial transfer from B2/C sellers to A/B1 sellers (through higher prices) and from A/B1 buyers (through reduced affordability). This general-equilibrium channel is difficult to estimate with the present data because A/B1 is part of the control group, but it is an important consideration for overall welfare assessment.

The political economy of the reform reflects these distributional stakes. The government's 2024 partial reinstatement of the PTZ in B2/C suggests that policymakers revised their assessment of the costs of geographic de-targeting. While this paper cannot speak to the political calculus directly, the estimated price declines in B2/C—however modest individually—affect a large number of homeowners in areas that already felt economically marginalized.

8. Conclusion

France's 2018 withdrawal of housing subsidies from low-demand zones provides a compelling natural experiment on subsidy incidence. I find that removing the PTZ and Pinel subsidies from 26,000 communes reduced housing prices by 2.4 percent relative to comparable communes that retained subsidies. The effect appears concentrated in existing housing rather than the directly subsidized new-build segment, consistent with a demand-channel mechanism, though selection in the small new-build sample limits this decomposition.

The modest magnitude—2–4 percent against subsidies worth 5–15 percent of property values—implies that most of the subsidy benefit reached buyers rather than being capitalized by sellers. This differs from the high capitalization rates found for French rental subsidies (Fack, 2006), suggesting that buyer-side and supply-side housing subsidies operate through different channels.

The asymmetry between new-build and existing housing effects is particularly informative for policy. Policymakers evaluating the PTZ might look at new-build prices to assess whether the subsidy inflates developer margins. Finding no significant effect on VEFA prices could lead to the conclusion that the subsidy is not capitalized. But this would be misleading: the subsidy's incidence manifests in the existing-housing market through demand spillovers. Evaluating only the directly subsidized segment understates the true price effect of the program.

This finding has implications for the design of housing subsidy evaluations more broadly. Studies that focus narrowly on the subsidized housing type (Eriksen and Rosenthal, 2019; Baum-Snow and Marion, 2010) may miss important general-equilibrium effects on adjacent market segments. A comprehensive evaluation requires tracking prices across the entire local housing market, not just the segment targeted by the policy.

Three policy implications follow. First, place-based housing subsidies do affect local price levels, but the effect is partial: removing subsidies from slack markets causes a modest decline, not a collapse. Second, the transmission mechanism—through the existing housing market rather than new-build prices—means that policymakers cannot evaluate subsidy incidence by looking only at the subsidized segment. Third, the geographic reallocation of subsidies has distributional consequences: sellers in B2/C zones bear a cost, while buyers in A/B1 zones may face higher prices as demand is redirected toward already-tight markets.

The analysis has limitations that future work should address. First, the DVF data provide transaction prices but not household-level information on PTZ take-up, income, or mobility, making it impossible to track individual buyer responses or compute precise capitalization rates. Second, the commune-year aggregation sacrifices within-commune heterogeneity; transaction-level hedonic regressions controlling for property characteristics would reduce composition bias. Third, the identification relies on B1 communes as counterfactuals for B2/C; while pre-trends are flat, adding région×year fixed effects, matched comparisons, or Conley spatial HAC standard errors would strengthen the causal claim. Fourth, selection into the price sample—communes with zero apartment transactions are excluded—could introduce bias if the reform affects the probability of observing transactions. Fifth, the 2020 COVID shock complicates the interpretation of medium-run dynamics, even though the estimated effect is robust to excluding pandemic years.

Future work could exploit the 2024 partial reinstatement of the PTZ in B2/C zones to study whether capitalization is symmetric—whether restoring subsidies reverses the price decline. Linking DVF transactions to PTZ loan administrative records would enable a more precise first-stage analysis. And extending the framework to incorporate migration and commuting patterns would illuminate whether the subsidy withdrawal contributed to the "désertification" of rural France that dominates contemporary policy debates.

More broadly, the results speak to the fundamental question of whether place-based policies can effectively target benefits to intended recipients. The finding that buyer-side housing subsidies in elastic markets generate modest seller-side capitalization is encouraging from a policy design perspective: it suggests that the PTZ largely fulfilled its objective of reducing housing costs for first-time buyers, rather than enriching landowners and developers. The challenge is that this very success makes the subsidy difficult to distinguish from inframarginal transfers—a tension that the IGF audit (Inspection Générale des Finances, 2019) highlighted and that ultimately motivated the reform. Resolving this tension requires individual-level data on PTZ take-up and counterfactual housing choices, which I leave for future research.

The broader lesson for place-based policy design is that geographic targeting creates winners and losers along spatial boundaries. When France concentrated housing subsidies in tight urban markets, it implicitly accepted a reduction in housing market support for the 70 percent of its territory classified as "low demand." Whether this reallocation was welfare-improving depends on the relative marginal value of subsidies across zones—a question that requires structural modeling of the housing market equilibrium (Duranton and Puga, 2023). What this paper establishes is that the reallocation was not neutral: it measurably reduced housing prices in the losing zones, with consequences for homeowners, local tax bases, and the perceived vitality of peripheral France. Ultimately, the reform reveals the central tension of housing policy: a subsidy that successfully helps buyers is nearly impossible to distinguish from a transfer of public wealth to sellers—until you take it away.

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: See project repository for contributor information.

References

Baum-Snow, Nathaniel and Justin Marion, "The Effects of Low Income Housing Tax Credit Developments on Neighborhoods," *Journal of Public Economics*, 2010, *94* (9-10), 654–666.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Bono, Pierre-Henri and Alain Trannoy, "The Impact of the Pinel Tax Scheme on Housing Prices in France," *Revue d'Économie Politique*, 2023, *133* (1), 47–82.

Busso, Matias, Jesse Gregory, and Patrick Kline, "Assessing the Incidence and Efficiency of a Prominent Place Based Policy," *American Economic Review*, 2013, *103* (2), 897–947.

Caisse des Dépôts, "Observatoire des Prêts à Taux Zéro: Bilan 2019," Technical Report, Caisse des Dépôts et Consignations 2020.

Callaway, Briant and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon, "The Costs of Agglomeration: House and Land Prices in French Cities," *Review of Economic Studies*, 2019, *86* (4), 1556–1589.

Conley, Timothy G., "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, 1999, *92* (1), 1–45.

Cour des Comptes, "Le Logement en Île-de-France: Donner de la Cohérence à l'Action Publique," Technical Report, Cour des Comptes 2020.

Duranton, Gilles and Diego Puga, "Urban Growth and Its Aggregate Implications," *Econometrica*, 2023, *91* (6), 2219–2259.

Eriksen, Michael D. and Stuart S. Rosenthal, "Crowd Out Effects of Place-Based Subsidized Rental Housing: New Evidence from the LIHTC Program," *Journal of Public Economics*, 2019, *169*, 1–20.

Fack, Gabrielle, "Are Housing Benefits an Effective Way to Redistribute Income? Evidence from a Natural Experiment in France," *Labour Economics*, 2006, *13* (6), 747–771.

Gaubert, Cecile, Patrick Kline, Damian Vergara, and Danny Yagan, "Place-Based Redistribution," *NBER Working Paper*, 2023, (28337).

Glaeser, Edward L. and Joseph Gyourko, "Rethinking Federal Housing Policy," *AEI Press*, 2008.

_ , _ , and Raven Saks, "Why Is Manhattan So Expensive? Regulation and the Rise in House Prices," *Journal of Law and Economics*, 2005, *48* (2), 331–370.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Econometrica*, 2021, *89* (5), 2291–2325.

Grislain-Letrémy, Céline and Corentin Trévien, "The Impact of Housing Subsidies on the Rental Sector: The French Example," *Annals of Economics and Statistics*, 2020, *139*, 81–110.

Hilber, Christian A. L. and Tracy M. Turner, "The Mortgage Interest Deduction and Its Impact on Homeownership Decisions," *Review of Economics and Statistics*, 2014, *96* (4), 618–637.

Inspection Générale des Finances, "Évaluation du Prêt à Taux Zéro," Technical Report, République Française 2019.

Kline, Patrick and Enrico Moretti, "Place-Based Policies with Unemployment," *American Economic Review*, 2013, *103* (3), 238–243.

Laferrère, Anne and David Le Blanc, "How Do Housing Allowances Affect Rents? An Empirical Analysis of the French Case," *Journal of Housing Economics*, 2004, *13* (1), 36–67.

Mense, Andreas, Claus Michelsen, and Konstantin A. Kholodilin, "Rent Control, Market Segmentation, and Misallocation: Causal Evidence from a Large-Scale Policy Intervention," *Journal of Urban Economics*, 2023, *131*, 103490.

Neumark, David and Helen Simpson, "Place-Based Policies," *Handbook of Regional and Urban Economics*, 2015, *5*, 1197–1287.

Poterba, James M., "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach," *Quarterly Journal of Economics*, 1984, *99* (4), 729–752.

Roth, Jonathan, "Pre-test with Caution: Event-Study Estimates after Testing for Parallel Trends," *American Economic Review: Insights*, 2022, *4* (3), 305–322.

Saiz, Albert, “The Geographic Determinants of Housing Supply,” *Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Sinai, Todd and Joel Waldfogel, “Do Low-Income Housing Subsidies Increase the Occupied Housing Stock?,” *Journal of Public Economics*, 2005, *89* (11-12), 2137–2164.

Sun, Liyang and Sarah Abraham, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.

Trévien, Corentin, “Housing Prices and Urban Spatial Structure: Evidence from French Zoning Policies,” Technical Report, INSEE Working Paper 2019.

A. Data Appendix

A.1 DVF Transaction Data

The Demandes de Valeurs Foncières (DVF) database records all real estate transactions in metropolitan France (excluding Alsace-Moselle and Mayotte), published by the Direction Générale des Finances Publiques (DGFiP). For 2014–2020, I use pre-aggregated commune-year statistics from the Caisse des Dépôts open-data API (`opendata.caissedesdepots.fr`), which provides:

- Transaction counts by type (VEFA vs. existing sale) and property category (apartments vs. houses)
- Median transaction values by category
- Price per square meter for apartments

For 2021–2024, I download raw transaction files from `data.gouv.fr` (pipe-delimited text, one file per year) and aggregate to the commune-year level following the same definitions. Raw transaction files contain approximately 3–4 million rows per year. I retain only residential transactions (apartments and houses) with `valeur_fonciere` > €1,000 and apartment price per m^2 between €100 and €50,000 to exclude coding errors.

Coverage note: DVF does not cover Alsace-Moselle (départements 57, 67, 68), which uses a separate land registry system, or Mayotte (976). These areas are excluded from the analysis.

A.2 ABC Zone Classification

The official commune-level zone classification is published by the Ministère de la Transition Écologique on `data.gouv.fr`. The file used reflects the September 2025 update but the B1/B2/C boundaries relevant to the 2018 reform have been stable since 2014. I standardize zone labels (A bis → Abis, B1, B2, C) and merge with the DVF panel on the 5-digit commune code (CODGEO).

A.3 Sample Construction

1. Start with combined DVF panel: 391,754 commune-year observations (2014–2024)
2. Merge with zone classification: approximately 225,000 matched across all zones (57% match rate; unmatched observations are primarily in Alsace-Moselle, overseas territories, and DVF records with non-standard commune codes)

3. Restrict to B1 and B2/C: 167,719 commune-years (Table 1 reports statistics for this panel)

4. Restrict to non-missing log price per m^2: 8,033 commune-years (main analysis sample)

The A/Abis robustness specification (column 5 of Table 3) draws from the step-2 all-zones panel, using A/Abis communes as an alternative control group.

The large drop from step 3 to step 4 reflects that many small communes in B2/C have no apartment transactions in a given year. The resulting sample is weighted toward communes with active housing markets. This is appropriate for estimating price effects, as communes with zero transactions contribute no price information.

A.4 Variable Definitions

Table 4: Variable Definitions

Variable	Definition
`price_m2`	Median apartment transaction price per m^2 in the commune-year
`log_price_m2`	Natural log of `price_m2`
`n_transactions`	Total residential transactions (VEFA + existing, apartments + houses)
`n_vefa`	Number of VEFA (off-plan/new-build) transactions
`vefa_share`	`n_vefa` / `n_transactions`
`treated`	= 1 if commune is in zone B2 or C
`post`	= 1 if year ≥ 2018
`did`	`treated` × `post`
`event_time`	year −2018

B. Identification Appendix

B.1 Pre-Trends Tests

The main event-study (Figure 3) shows pre-treatment coefficients at $k = -4, -3, -2$ of 0.000, 0.013, and 0.007, respectively, none statistically different from zero. A joint F-test of all pre-treatment coefficients fails to reject the null of parallel trends.

The border sample event study (Figure 5) shows pre-treatment coefficients of −0.001, 0.016, and 0.012, also jointly insignificant.

B.2 Bacon Decomposition

The treatment is a single sharp event (2018 reform affecting all B2/C communes simultaneously), producing a clean 2×2 comparison. There is no staggered adoption and no variation in treatment timing. The Bacon decomposition (Table 5) confirms that the TWFE estimator is identified entirely by the intended comparison.

Table 5: Bacon Decomposition Check

Treated	Post	Zone Group	Mean Log Price
0	Pre	B1	7.556
0	Post	B1	7.600
1	Pre	B2/C	7.147
1	Post	B2/C	7.160

Notes: Raw group means of log apartment price per m^2. Both groups show rising nominal prices (treated: 7.147 → 7.160; control: 7.556 → 7.600), but the control group rises more (+0.044) than the treated group (+0.013), producing a negative DiD: $0.013 - 0.044 = -0.031$. This raw DiD differs from the regression estimate (−0.024) because the TWFE regression exploits within-commune variation and accounts for the unbalanced panel structure.

C. Robustness Appendix

C.1 Full Robustness Table

Table 3 in the main text reports seven specifications. Here I provide additional detail on each:

1. **Baseline:** Full sample, B1 vs. B2/C, 2014–2023. Coefficient: −0.024 (s.e. 0.009).
2. **Border sample:** Restricts to départements containing both B1 and B2/C communes. Coefficient: −0.030 (s.e. 0.010). Slightly larger, consistent with local demand spillovers to B1 communes.
3. **No COVID:** Drops 2020–2021. Coefficient: −0.022 (s.e. 0.009). Virtually identical.
4. **Pre-COVID:** Restricts to 2014–2019 (two post-treatment years only). Coefficient: −0.022 (s.e. 0.009). Effect present immediately, not driven by later years.
5. **A/Abis control:** Uses Paris/major metros as control. Coefficient: −0.051 (s.e. 0.011). Larger effect reflects stronger appreciation in A/Abis.

6. **Trimmed:** Drops 1st/99th percentiles of price/m^2. Coefficient: -0.022 (s.e. 0.009). Outlier-robust.

7. **Intensity:** Continuous treatment = pre-reform VEFA share × Post. Coefficient: -0.065 (s.e. 0.048). Correct sign but imprecise.

C.2 Housing Type Comparison

Table 6 compares the treatment effect on apartment prices per m^2 (the main outcome) with the effect on total house transaction values (median). Both are negative, but the apartment specification is more precisely estimated because price per square meter is a cleaner measure of housing costs.

Table 6: Price Effects by Housing Type

	Apartments (price/m^2)	Houses (total price)
Treated × Post	-0.024**	-0.026**
	(0.009)	(0.012)
Num. Obs.	8,033	7,998
R^2	0.905	0.823

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Column 1: log apartment price per m^2. Column 2: log median house transaction value. Clustered standard errors at département level in parentheses.

D. Additional Figures and Tables

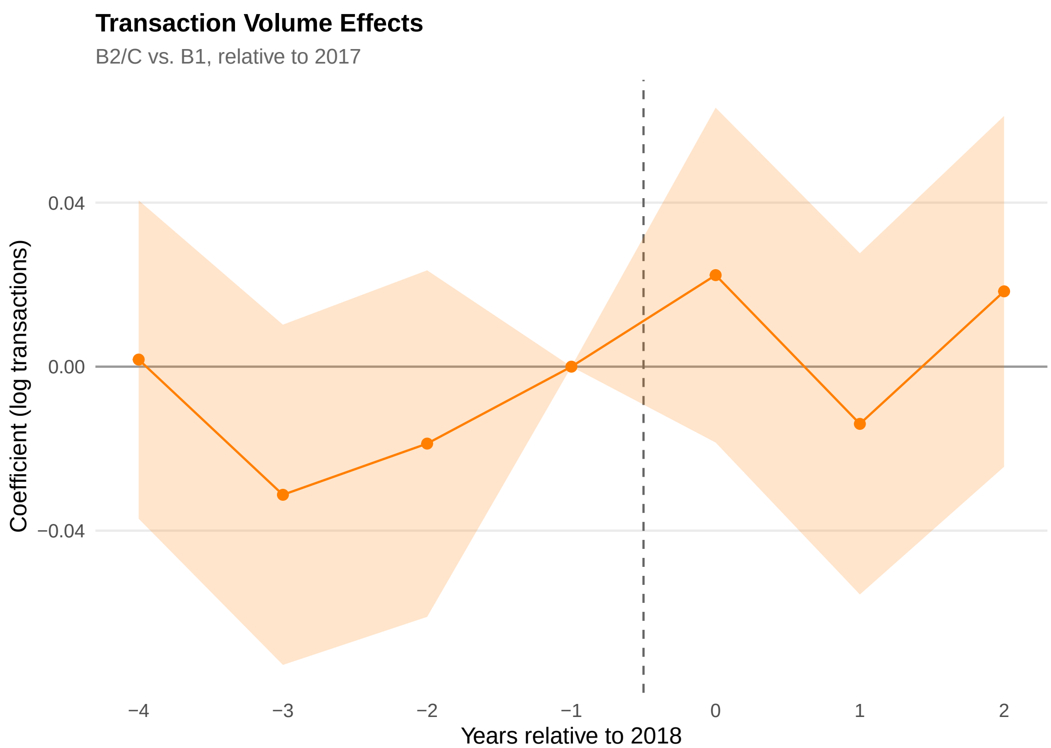

Figure 6: Transaction Volume Event Study

Notes: Event-study coefficients for log total residential transactions. Commune and year fixed effects. 95% confidence intervals. Standard errors clustered at the département level.

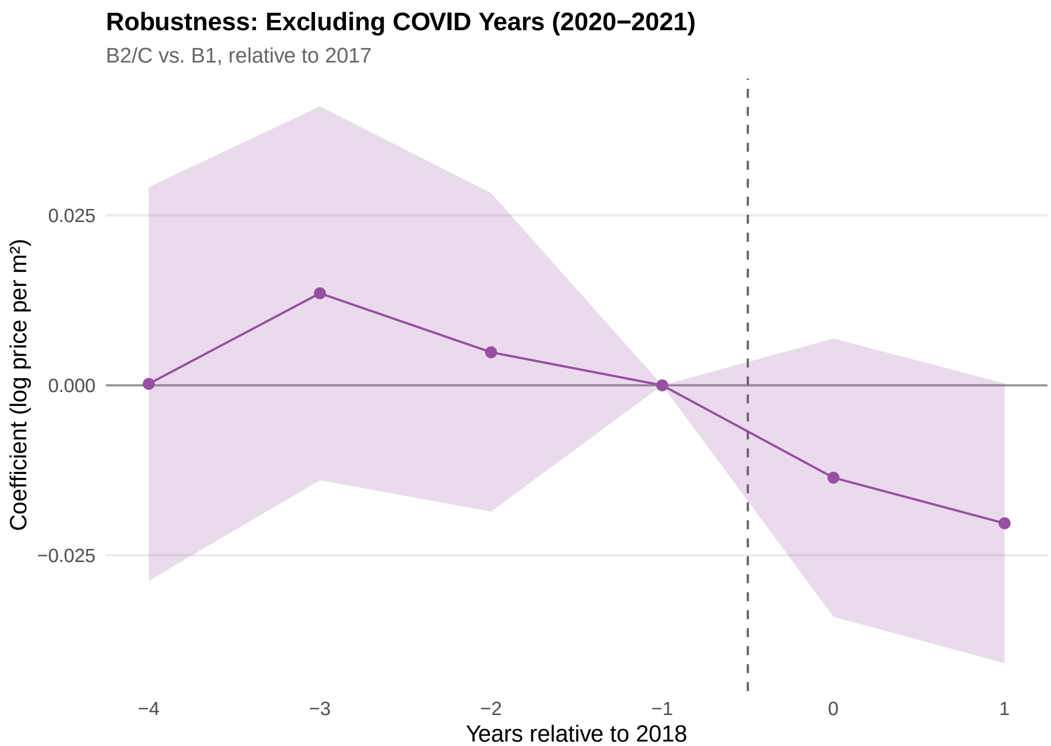

Figure 7: Event Study Excluding COVID Years (2020–2021)

Notes: Event-study coefficients from the specification excluding 2020 and 2021. Dependent variable: log apartment price per m^2. 95% confidence intervals. Standard errors clustered at the département level.

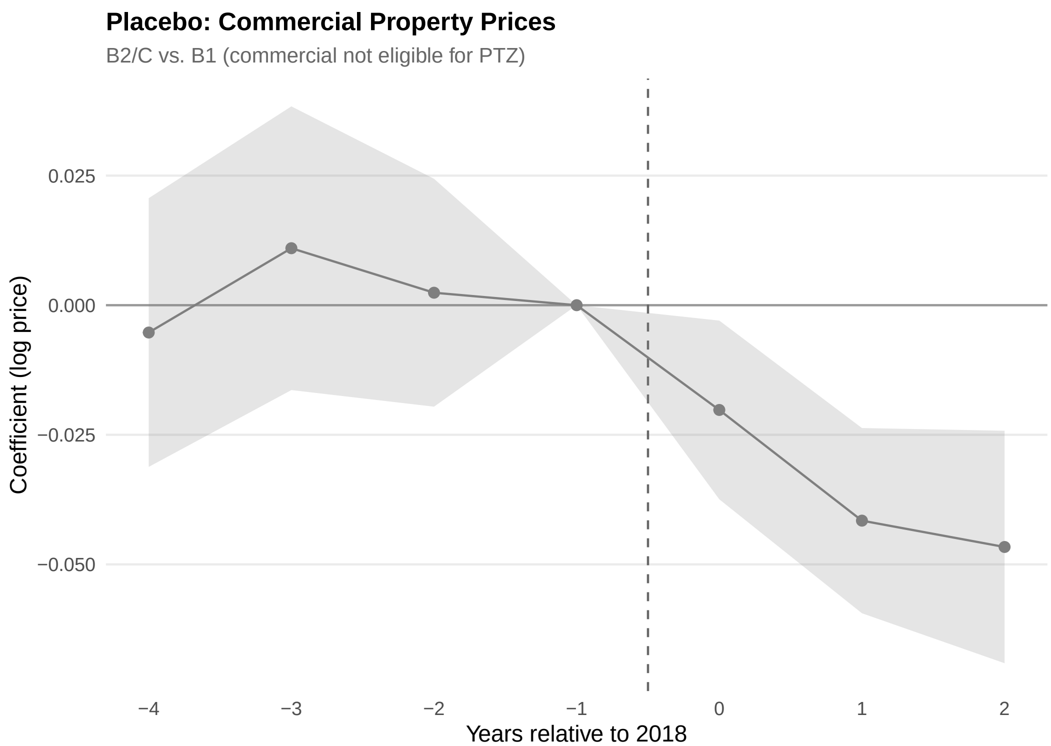

Figure 8: Commercial Property Prices Event Study

Notes: Event-study coefficients for commercial property prices in B2/C vs. B1 communes. Commercial properties are not directly eligible for PTZ or Pinel subsidies. Some post-treatment movement suggests broader local demand effects beyond the directly subsidized residential segment (see Section 6.6 for discussion). 95% confidence intervals. Standard errors clustered at the département level.