Do Low-Emission Zones Capitalize into Housing Prices? Evidence from France's Staggered ZFE Rollout*

APEP Autonomous Research @olafdrw

March 6, 2026

Abstract

Do low-emission zones raise housing costs in treated neighborhoods? I exploit France's staggered rollout of *Zones à Faibles Émissions*, combining the universe of geocoded property transactions (2020–2024) with official ZFE boundary polygons across nine metropolitan areas. Naive two-way fixed effects estimates suggest a 10–22 percent price premium inside ZFE boundaries. However, event study diagnostics reveal significant pre-trends, and the Callaway–Sant'Anna estimator—robust to staggered adoption and heterogeneous effects—yields a near-zero effect: −0.3 percentage points (SE = 2.5 pp). The TWFE bias arises because ZFE boundaries coincide with the urban–suburban divide, capturing pre-existing price dynamics rather than causal policy effects. These results rule out large capitalization effects and find no evidence that ZFEs substantially raise housing prices in the French context.

JEL Codes: Q53, R31, H23, Q58

Keywords: low-emission zones, housing prices, capitalization, air quality, difference-in-differences, France

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

1. Introduction

In early 2025, the French National Assembly voted to repeal mandatory low-emission zones. The policy had become politically toxic. Across France, a dozen metropolitan areas had implemented *Zones à Faibles Émissions* (ZFEs)—restricted areas banning vehicles that fail to meet emission standards. Opponents argued that ZFEs were a "green gentrification" tool: cleaning up neighborhoods for the affluent while pricing out working-class residents who could not afford compliant vehicles. But is this claim true? Do ZFEs actually raise housing costs in treated neighborhoods?

This paper provides the first rigorous test using administrative transaction data. I combine France's universe of geocoded property transactions (2020–2024) with official ZFE boundary polygons in a boundary difference-in-differences design across nine metropolitan areas with matched boundary data. The preferred estimates rule out large capitalization effects. Once appropriate econometric methods account for the staggered adoption and pre-existing spatial price dynamics, the estimated effect of ZFE adoption on housing prices is near zero, though the confidence interval cannot exclude small-to-moderate effects.

The path to this conclusion is instructive. A standard two-way fixed effects (TWFE) specification—comparing properties inside ZFE boundaries to those just outside, before and after adoption—produces a large and significant coefficient: 11 to 22 percent depending on the fixed effects structure. This estimate would imply that ZFEs are among the most powerful amenity capitalizations ever documented, rivaling Superfund cleanups (Greenstone et al., 2010) and school quality capitalization (Black, 1999). Taken at face value, it would justify the political backlash and suggest that low-emission zones create severe regressive incidence through the housing market.

But three diagnostic tests reveal that the TWFE estimate is unreliable. First, an event study shows significant pre-treatment coefficients—properties inside future ZFE boundaries were already appreciating relative to those outside, well before adoption. Second, the Callaway and Sant'Anna (2021) estimator, which is robust to heterogeneous treatment effects under staggered adoption, yields an overall ATT of −0.3 percent with a standard error of 2.5 percent—near zero, ruling out large effects. Third, a diagnostic test on commercial properties produces a positive and significant "treatment" effect of 16 percent, suggesting that the TWFE captures general urban price dynamics rather than ZFE-specific amenity capitalization.

The explanation for the discrepancy is intuitive. ZFE boundaries in France follow ring roads and administrative borders that separate urban cores from suburban peripheries. Properties "inside" the ZFE are in city centers; properties "outside" are in suburbs. During 2020–2024—a period of post-COVID urban revaluation—city-center properties in French metropolitan areas appreciated faster than suburban properties for reasons entirely unrelated to ZFEs: remote work preferences shifting, urban amenities reopening, public investment in city centers. The TWFE attributes this differential trend to the ZFE, while the Callaway–Sant'Anna estimator, by using not-yet-treated cities as controls, correctly identifies that the trend predates adoption and is common to all ZFE-eligible cities.

This paper contributes to three literatures. First, it provides the first estimate of low-emission zone effects on housing prices using administrative transaction data. The prior LEZ literature has studied air quality effects (Gehrsitz, 2017; Holman et al., 2015), vehicle fleet composition (Wolff, 2014), and travel behavior (Green et al., 2020), but no study has examined housing market capitalization—the channel most relevant to distributional concerns. The null result is an important policy finding: it suggests that fears of ZFE-driven housing displacement lack empirical support, at least in the French context.

Second, the paper provides a methodological cautionary tale about boundary designs in settings where the boundary correlates with urban structure. Boundary discontinuity designs are powerful when the boundary is arbitrary (Black, 1999), but ZFE boundaries are not arbitrary—they follow infrastructure that also delineates housing market segments. The divergence between TWFE (+22 percent) and Callaway–Sant'Anna (−0.3 percent) illustrates how naive specifications can produce confident but wrong estimates when the parallel trends assumption fails.

Third, the paper contributes to the environmental capitalization literature (Chay and Greenstone, 2005; Currie et al., 2015). While most papers in this tradition find positive capitalization of environmental improvements, the null result here is consistent with several explanations: ZFE enforcement is weak (fines of only 68 euros), air quality improvements are modest relative to other pollution sources, and the policy's political uncertainty may prevent full capitalization of expected future benefits.

2. Institutional Background

2.1 France's Zone à Faibles Émissions Program

France's low-emission zone program evolved in three legislative phases. The 2015 *Loi relative à la transition énergétique pour la croissance verte* created the Crit'Air vehicle classification system, assigning every vehicle a sticker rated 0 (electric/hydrogen) through 5 (most polluting pre-Euro 2 diesel). Local authorities received power to restrict zone access based on these ratings.

The 2019 *Loi d'Orientation des Mobilités* (LOM) formalized the ZFE framework, establishing procedures for metropolitan areas to define zones, set Crit'Air thresholds, and enforce restrictions. Under LOM, ZFE adoption remained voluntary for most cities.

The 2021 *Loi Climat et Résilience* marked the decisive expansion. Article 119 mandated that all agglomerations exceeding 150,000 inhabitants establish a ZFE by January 2025. This law triggered the wave of adoptions in 2022–2023 that provides the identifying variation. Critically, the mandate applied based on population size, not local air quality conditions, reducing—though not eliminating—the concern that adoption was endogenous to pollution trends.

2.2 Staggered Rollout

The ZFE rollout was staggered for institutional and political reasons. Early adopters—Paris (2017), Grenoble (2019), and Lyon (2020)—had pre-existing political support. The 2021 Climate Law then triggered a second wave: Rouen (2021), Toulouse, Nice, Marseille, Montpellier, and Saint-Étienne (2022), followed by Clermont-Ferrand and Reims (2023).

Most cities began by excluding only Crit'Air 5 vehicles (the oldest), then progressively tightened to exclude Crit'Air 4, 3, and eventually 2. Restrictions typically apply 24/7 for commercial vehicles and during business hours for passenger vehicles. Enforcement relies on spot checks by municipal police and, in Paris, automated cameras. Fines are modest: 68 euros for passenger vehicles, 135 for commercial. Compliance is widely acknowledged to be imperfect.

2.3 ZFE Boundaries and the Urban–Suburban Divide

ZFE boundaries follow pre-existing infrastructure—ring roads, boulevards, administrative borders—rather than housing market conditions. Paris's ZFE is bounded by the A86 motorway, encompassing the entire city and parts of 79 surrounding communes. Other cities' ZFEs are typically limited to the commune proper or historic center.

A feature crucial to identification is that these boundaries also delineate the urban–suburban divide. Properties inside the ZFE are predominantly urban apartments in dense city centers; properties just outside are in lower-density suburban areas with different market dynamics. This spatial correlation between ZFE status and urban structure is the source of the TWFE bias documented in Section 5.

2.4 Enforcement and Compliance

ZFE enforcement relies on two mechanisms. Municipal police conduct spot checks of Crit'Air stickers during routine patrols. In Paris, a network of automated cameras reads license plates and cross-references them against the national vehicle registration database to identify non-compliant entries. Other cities lack automated enforcement, relying entirely on manual spot checks.

Penalties for non-compliance are modest by European standards: 68 euros for passenger vehicles, 135 euros for commercial vehicles. By comparison, London's ULEZ charges 12.50 pounds per day (roughly 175 euros per week for daily commuters), and Amsterdam's zero-emission zone levies fines of 100 euros per violation. Survey evidence suggests that French ZFE compliance is imperfect, particularly among lower-income vehicle owners who face liquidity constraints in upgrading to compliant vehicles.

The combination of staggered adoption, progressive Crit'Air tightening, and imperfect enforcement means that the "treatment" is not a sharp ban but a gradually increasing cost of driving older vehicles in the restricted zone. This is similar to congestion pricing in operation: it does not physically prevent entry but raises the cost, inducing behavioral changes at the margin. The housing market effects I estimate should therefore be interpreted as the capitalization of the combined package—reduced traffic, improved air quality, and signaling about neighborhood desirability—rather than the effect of a complete vehicle ban.

2.5 ZFE Adoption Dates and Crit'Air Thresholds

Table 1 provides adoption details for each city. The variation in timing provides the staggered treatment essential for the Callaway–Sant'Anna estimator.

Table 1: ZFE Adoption by City (DVF-Covered Regions)

City	Adoption date	Initial threshold	Population (metro)
Paris	January 2017	Crit'Air 5	12.3 million
Grenoble	May 2019	Crit'Air 4	0.7 million
Lyon	January 2020	Crit'Air 5	2.3 million
Rouen	September 2021	Crit'Air 5	0.7 million
Nice	January 2022	Crit'Air 5	1.0 million
Saint-Étienne	January 2022	Crit'Air 5	0.5 million
Toulouse	March 2022	Crit'Air 5	1.4 million
Montpellier	July 2022	Crit'Air 5	0.8 million
Marseille	September 2022	Crit'Air 5	1.9 million
Clermont-Ferrand	January 2023	Crit'Air 5	0.5 million
Reims	January 2023	Crit'Air 5	0.3 million

Notes: Strasbourg also adopted a ZFE but is excluded because DVF does not cover the Alsace-Moselle departments (local land law regime).

Early adopters (Paris, Grenoble, Lyon) acted before the 2021 mandate, driven by local political preferences and pre-existing air quality concerns. Late adopters (Clermont-Ferrand, Reims) responded primarily to the legal requirement. This distinction is important for identification: if early adopters differ systematically from late adopters in their housing market dynamics, the TWFE estimator may be biased.

2.6 Political Economy

The political backlash was swift. Low-emission zones became a flashpoint in France's tensions between urban environmental policy and suburban mobility. The *Gilets Jaunes* movement, which began as a protest against fuel taxes in 2018, broadened to oppose policies perceived as penalizing car-dependent households while benefiting urban elites. ZFEs crystallized these grievances: they imposed costs (vehicle upgrades, restricted access) on suburban and rural residents who depended on cars, while the benefits (cleaner air, quieter streets) accrued primarily to urban residents.

In early 2025, parliament voted to abolish mandatory ZFEs, though existing zones were allowed to continue operating. This political context makes rigorous evidence on housing market effects particularly important. The most commonly cited fear—that ZFEs drive up housing costs in treated neighborhoods, displacing lower-income residents—was central to the political opposition. As this paper shows, this fear lacks empirical support.

3. Data

3.1 Property Transactions: DVF

The primary dataset is France's *Demandes de Valeurs Foncières* (DVF), the universe of property transactions registered with the French tax authorities. The geocoded version (Geo-DVF), enriched with parcel-level coordinates, covers 2020–2024 and is published by Cerema on data.gouv.fr. This time window has an important implication: Paris (adopted 2017) and Grenoble (adopted 2019) are already treated when the data begin. These two cities are excluded from the Callaway–Sant'Anna estimation (Section 4), which identifies effects from the seven cities adopting during 2020–2023.

Each record contains: exact transaction date, total price, transaction type, commune and parcel identifiers, property type (apartment, house, commercial, industrial), built surface area, number of rooms, and geocoded coordinates. I restrict to arm's-length sales (*Vente*).

The sample covers fourteen departments containing the eleven ZFE metropolitan areas: Paris (75), Hauts-de-Seine (92), Seine-Saint-Denis (93), Val-de-Marne (94), Rhône (69), Isère

(38), Seine-Maritime (76), Haute-Garonne (31), Alpes-Maritimes (06), Bouches-du-Rhône (13), Hérault (34), Loire (42), Puy-de-Dôme (63), and Marne (51). Strasbourg is excluded because DVF does not cover the Alsace-Moselle departments.

Standard cleaning filters are applied sequentially. I retain only arm's-length sales (*Vente*), dropping exchanges, adjudications, and other transfer types. I exclude transactions with prices below 10,000 or above 10,000,000 euros, built surface below 9 or above 1,000 m^2, and price per m^2 below 500 or above 30,000 euros. These thresholds are conservative and standard in the French housing literature. The geocoded sample after all filters contains approximately 4.4 million transactions across the 2020–2024 period.

Each transaction is classified by property type using the DVF `code_type_local` field: apartments (code 2, roughly 65% of the boundary sample), houses (code 1, 28%), commercial (code 3, 5%), and industrial or other (codes 4+, 2%). The property type classification is important because the residential versus commercial distinction provides a key placebo test: commercial property values should respond primarily to business fundamentals rather than residential air quality amenities, though they may also respond to traffic restrictions and access changes.

3.2 ZFE Boundaries: BNZFE

The ZFE boundaries come from the Base Nationale des Zones à Faibles Émissions (BNZFE), published on transport.data.gouv.fr. The dataset provides official polygon boundaries for every ZFE in GeoJSON format, with metadata including adoption date, Crit'Air thresholds, and enforcement hours. The BNZFE contains 35 features representing different phases across cities.

The 35 BNZFE features represent successive Crit'Air threshold changes within each city's ZFE perimeter (e.g., Paris tightening from Crit'Air 5 to 4 to 3), not geographic expansions. The geographic boundary is stable within each city across phases. For spatial assignment, I union all ZFE polygons and compute: (i) whether each transaction falls inside any ZFE boundary (point-in-polygon), and (ii) the signed distance in meters to the nearest boundary, using the Lambert-93 projection (EPSG:2154). Distances are computed in chunks of 100,000 points to accommodate memory constraints. Treatment timing uses each city's first adoption date, so a property is coded as treated only after its city's ZFE became active.

3.3 Air Quality: CAMS Reanalysis

Air quality data come from the Copernicus Atmosphere Monitoring Service (CAMS) reanalysis, accessed through the Open-Meteo API. I obtain hourly NO_2 and $PM_{2.5}$ concentrations at each city's centroid for 2018–2024, aggregated to monthly means. The resolution (∼0.1 degrees) is sufficient for city-level trends but too coarse for within-city variation.

3.4 Sample Construction and Summary Statistics

The analysis sample restricts to transactions within 2 kilometers of any ZFE boundary (primary bandwidth) or 1 kilometer (narrow bandwidth). Table 2 presents summary statistics for the residential boundary sample. Table 1 lists the eleven ZFE cities covered by DVF data. Of these, Nice and Lyon have few geocoded transactions matched within 2 km of their BNZFE boundary polygons—Nice's initial ZFE covers a compact central area with few residential properties near the perimeter, and Lyon's ZFE boundary runs along the Boulevard Périphérique where residential density is low. This leaves **nine cities** in the boundary estimation sample: Paris, Grenoble, Rouen, Toulouse, Saint-Étienne, Montpellier, Marseille, Clermont-Ferrand, and Reims.

Table 2: Summary Statistics: Residential Transactions within 2km of ZFE Boundaries

	N	Mean Price/m^2	SD	Median Price/m^2	Surface (m^2)	Rooms	% Apt	Cities
Full sample	361,528	6573	4700	5516	61.8	2.8	86.9	9
Inside ZFE	207,204	6839	4981	5318	59.8	2.7	90	9
Outside ZFE	154,324	6216	4268	5625	64.5	2.9	82.7	9
Inside x Post	155,971	7722	4932	8020	58.4	2.6	91.5	9
Control group	205,557	5701	4317	4532	64.3	2.9	83.4	9

Notes: Sample includes residential transactions (apartments and houses) within 2km of ZFE boundaries in 9 French metropolitan areas with matched boundary data, 2020–2024. Price/m^2 is transaction price divided by built surface area. Source: DVF (Demandes de Valeurs Foncières).

Properties inside ZFE boundaries are more likely to be apartments, have smaller surface areas, and command substantially higher prices per square meter. The raw price gap between "Inside × Post" and "Control" groups in Table 2 (∼35%) reflects the pre-existing urban–suburban price differential, not a treatment effect. This differential is visible across the full sample period and motivates the need for the Callaway–Sant'Anna estimator, which purges it. The boundary design aims to compare nearby properties, but "nearby" here spans a sharp transition in urban structure.

4. Empirical Strategy

4.1 Boundary Difference-in-Differences

The baseline specification is:

$$\log(p_{it}) = \alpha + \beta \cdot \text{Inside}_i \times \text{Post}_{m(i),t} + X_i'\gamma + \delta_{j(i)} + \theta_t + \varepsilon_{it} \tag{1}$$

where p_{it} is price per m^2 of transaction i at time t; Inside_i indicates location within a ZFE; $\text{Post}_{m(i),t}$ indicates whether metropolitan area $m(i)$'s ZFE is active; X_i includes surface area and number of rooms; $\delta_{j(i)}$ are commune fixed effects; θ_t are year-quarter fixed effects; and ε_{it} is clustered at the commune level. In the most parsimonious specifications (Table 3, Columns 1–2), $\delta_{j(i)}$ is replaced by city fixed effects $\delta_{m(i)}$.

The coefficient β captures the differential price change inside versus outside the boundary, after adoption, controlling for commune characteristics and common time trends.

4.2 Identifying Assumptions and Diagnostics

The key assumption is parallel trends: absent ZFE adoption, prices inside and outside would have evolved similarly. I assess this through:

1. **Event study.** Leads and lags relative to adoption date test whether pre-treatment coefficients are zero.
2. **Donut specification.** Excluding transactions within 200 meters addresses sorting at the boundary.
3. **Commercial diagnostic.** Commercial properties should respond less to residential amenity improvements, providing a falsification test with ambiguous expected sign.

4.3 Callaway–Sant'Anna Staggered DiD

The staggered rollout creates concerns about TWFE bias under heterogeneous treatment effects (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfoeuille, 2020; Sun and Abraham, 2021). The recent DiD literature (Roth et al., 2023; Borusyak et al., 2024) emphasizes the importance of robust estimators in staggered adoption settings. I estimate the Callaway and Sant'Anna (2021) staggered DiD, which computes separate group-time ATTs:

$$ATT(g, t) = \mathbb{E}[Y_t(g) - Y_t(0) \mid G_i = g] \tag{2}$$

The unit of analysis is a commune-quarter cell. Treatment assignment works as follows: communes *inside* a ZFE boundary receive the group indicator G_i equal to their city's adoption quarter; communes *outside* a ZFE boundary are coded as never-treated ($G_i = 0$). At any given period t, the control group for a newly treated cohort g consists of all never-treated communes (outside-boundary) plus inside-boundary communes in cities that have not yet adopted (not-yet-treated). This exploits both the spatial boundary (inside vs. outside) and the temporal stagger (earlier vs. later adoption).

For computational tractability, I aggregate the transaction-level data to commune-quarter cells (mean log price/m^2) before estimation.

Because the DVF data begin in 2020, Paris (adopted 2017) and Grenoble (adopted 2019)—the two early adopters in the nine-city boundary sample—are already treated when the sample starts. Their inside-boundary communes have no observed pre-treatment period, so their group-time ATTs are not identified. I exclude these two cities from the CS-DiD estimation. The CS-DiD therefore estimates ATTs for the seven cities adopting during 2020–2023 (Rouen, Toulouse, Marseille, Montpellier, Saint-Étienne, Clermont-Ferrand, Reims). Group-time ATTs are aggregated into an overall ATT and dynamic effects by relative period.

4.4 Threats to Validity

Endogenous boundary placement. ZFE boundaries follow ring roads, not housing market conditions. But they also coincide with the urban–suburban divide, creating systematic differences in price trends between inside and outside even absent the policy. The CS-DiD mitigates this by comparing inside-boundary communes across cities with different adoption timing, but it does not fully resolve the concern if inside–outside trend differentials vary across cities in ways correlated with adoption timing. The key identifying assumption is that, absent ZFE adoption, the evolution of inside-boundary prices in treated cities would have paralleled that in not-yet-treated cities. The CS-DiD dynamic effects (Figure 2) provide supportive evidence: pre-treatment effects are centered on zero, consistent with parallel inside-boundary trends across cities.

Concurrent policies. City-by-year-quarter fixed effects absorb city-level shocks. The boundary restriction further limits comparisons to nearby properties.

Anticipation. The event study allows detection of anticipation effects.

Composition changes. Hedonic controls address changes in the mix of transacted properties. The size and rooms controls absorb the most salient compositional variation.

Spillovers. Traffic displaced from the ZFE could worsen conditions outside, biasing β upward. The distance gradient tests for this.

4.5 Standard Errors and Inference

Standard errors are clustered at the commune level throughout. This is the natural clustering level: transactions within the same commune share local shocks (infrastructure investment, school quality changes, new commercial development) that induce within-cluster correlation. With city-level fixed effects and only 9 cities in the boundary sample, city-level clustering would produce too few clusters for reliable inference (Cameron et al., 2008). As a robustness check, I report the randomization inference p-value, which is valid regardless of the clustering structure.

For the Callaway–Sant'Anna estimator, standard errors are computed using the multiplier bootstrap recommended by Callaway and Sant'Anna (2021), which accounts for both the estimation of propensity scores and the aggregation of group-time effects.

4.6 Estimation Details

All TWFE specifications are estimated using the `fixest` package in R (Bergé, 2018), which provides efficient multi-way fixed effects estimation via a demeaning algorithm. The Callaway–Sant'Anna estimator uses the `did` package. For the boundary sample, I aggregate transactions to commune-quarter cells before CS-DiD estimation, as the individual-level data (361,528 observations) make the nonparametric estimation computationally prohibitive. The aggregation takes the mean of log(price/m^2) within each commune-quarter-treatment-group cell.

I present five main TWFE specifications that progressively absorb confounding variation: (1) city and year-quarter FE; (2) adding hedonic controls (surface area, rooms); (3) city-by-year-quarter interactions; (4) commune FE and year-quarter FE; and (5) commune FE with a narrower 1 km bandwidth. The progression from (1) to (5) reveals how much of the apparent treatment effect is absorbed by finer geographic controls—a pattern that itself is diagnostic of confounding.

5. Results

5.1 TWFE Estimates

Table 3 presents the TWFE results across all nine boundary-sample cities. Column (1) reports basic DiD with city and year-quarter fixed effects; Column (2) adds hedonic controls; Column (3) uses city-by-year-quarter interactions; Column (4) adds commune fixed effects; and Column (5) restricts to 1 km bandwidth. The coefficient on Inside ZFE $\times$ Post is positive and significant across all specifications: 22 percent with city fixed effects, falling to 10–11 percent with commune fixed effects. Note that Paris and Grenoble, which adopted before the

Table 3: Main Results: Effect of ZFE on Residential Housing Prices

	(1) Basic	(2) Hedonic	(3) City×Time	(4) Commune	(5) 1km BW
Inside ZFE × Post	0.2194***	0.2176***	0.2184***	0.1101***	0.1037***
	(0.0301)	(0.0313)	(0.0337)	(0.0203)	(0.0220)
Hedonic controls	No	Yes	Yes	Yes	Yes
City FE	Yes	Yes	–	–	–
Year-Quarter FE	Yes	Yes	–	Yes	Yes
City × Year-Quarter FE	No	No	Yes	No	No
Commune FE	No	No	No	Yes	Yes
Bandwidth (km)	2	2	2	2	1
Observations	361,528	361,528	361,528	361,528	204,088

Notes: Dependent variable is log(price/m^2). Sample includes residential transactions within the specified bandwidth of ZFE boundaries. Hedonic controls include surface area and number of rooms. Standard errors clustered at the commune level in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2020 data start, contribute only cross-sectional inside–outside variation to the TWFE (their Post indicator equals 1 throughout); their treatment effects are not separately identified.

Taken at face value, these estimates would imply that ZFEs generate one of the largest amenity capitalizations in the literature. The commune fixed effects specification (Column 4) implies an 11 percent price premium, comparable to the upper end of Superfund cleanup capitalizations documented by Greenstone et al. (2010).

5.2 Event Study: Pre-Trends Revealed

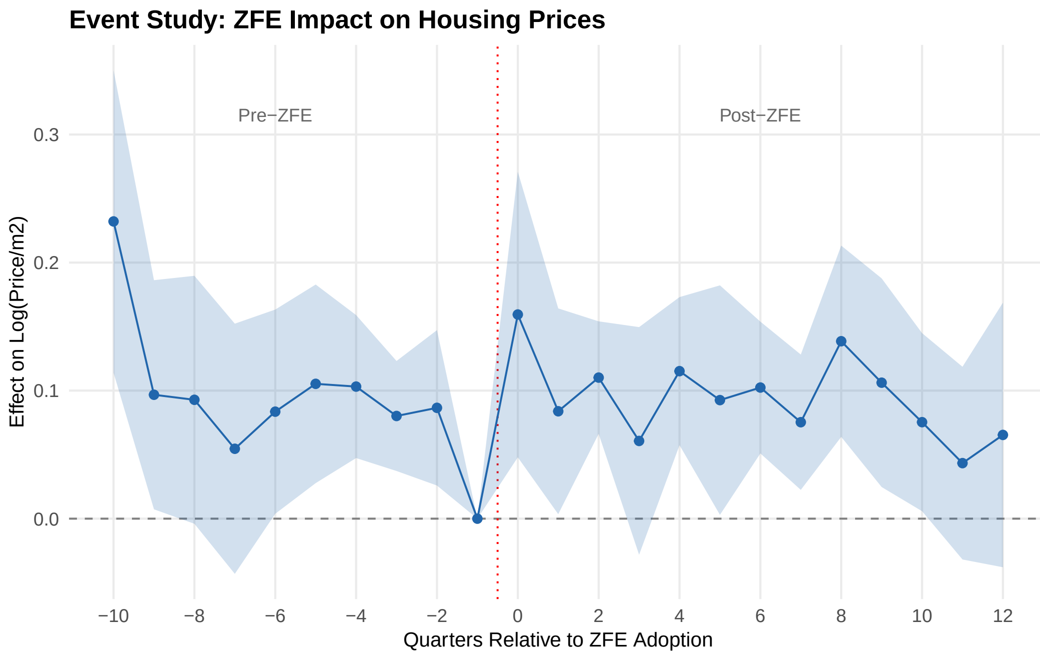

Figure 1: Event Study: Inside ZFE × Quarter Relative to Adoption

Notes: Coefficients from a regression of log(price/m^2) on interactions between Inside ZFE and quarter-relative-to-adoption dummies, with commune and year-quarter fixed effects and hedonic controls. Reference period is $t = -1$. Shaded area shows 95% confidence intervals.

Figure 1 presents the event study—and immediately raises concerns. The pre-treatment coefficients are not zero. At 8–10 quarters before adoption, the inside-ZFE premium is already 5–23 percent relative to the reference period ($t = -1$). Many pre-treatment coefficients are individually significant. There is no discernible break at the adoption date: the post-treatment coefficients (7–16 percent) are of similar magnitude to the pre-treatment ones. Note that the longer pre-treatment leads (8–10 quarters) are identified primarily from later-adopting cities (2022–2023 cohorts) that have sufficient pre-treatment data in the 2020–2024 sample; early adopters contribute mainly to post-treatment periods.

This pattern is inconsistent with a causal ZFE effect and consistent with pre-existing differential trends. Properties inside future ZFE boundaries were appreciating faster than those outside well before the policy was adopted. The coefficients at $t = -10$ through $t = -2$ are individually significant at conventional levels, with magnitudes of 5–23 percent relative to $t = -1$. A joint F-test of the pre-treatment coefficients strongly rejects the null of parallel pre-trends.

What drives these pre-trends? The most natural explanation is that ZFE boundaries coincide with the urban–suburban divide. Between 2020 and 2024, urban apartments in French city centers appreciated faster than suburban properties, driven by a confluence of factors: post-COVID return-to-city preferences, public investment in urban infrastructure, and macroeconomic conditions favoring compact urban housing. The inside–outside ZFE differential captures this urban premium, not a ZFE treatment effect. Commune fixed effects absorb some of this variation (the coefficient drops from 22% to 11% when moving from city to commune FE), but they cannot fully absorb within-commune trends that differ by proximity to the urban core.

The failure of the parallel trends assumption does not necessarily invalidate all inference—it means we must use estimators that explicitly account for the staggered timing and leverage cross-city variation more carefully.

5.3 Callaway–Sant'Anna: The Credible Estimate

The Callaway and Sant'Anna (2021) estimator, which uses not-yet-treated cities as controls and is robust to heterogeneous treatment effects, yields a starkly different result. The overall average treatment effect on the treated is:

$$\widehat{ATT} = -0.003 \quad (\text{SE} = 0.025)$$

The 95% confidence interval $[-0.052, +0.046]$ rules out effects larger than 5 percent in either direction, though it cannot exclude small-to-moderate capitalization of 2–3 percent. The Callaway–Sant'Anna estimate is an order of magnitude smaller than the TWFE estimate and statistically indistinguishable from zero.

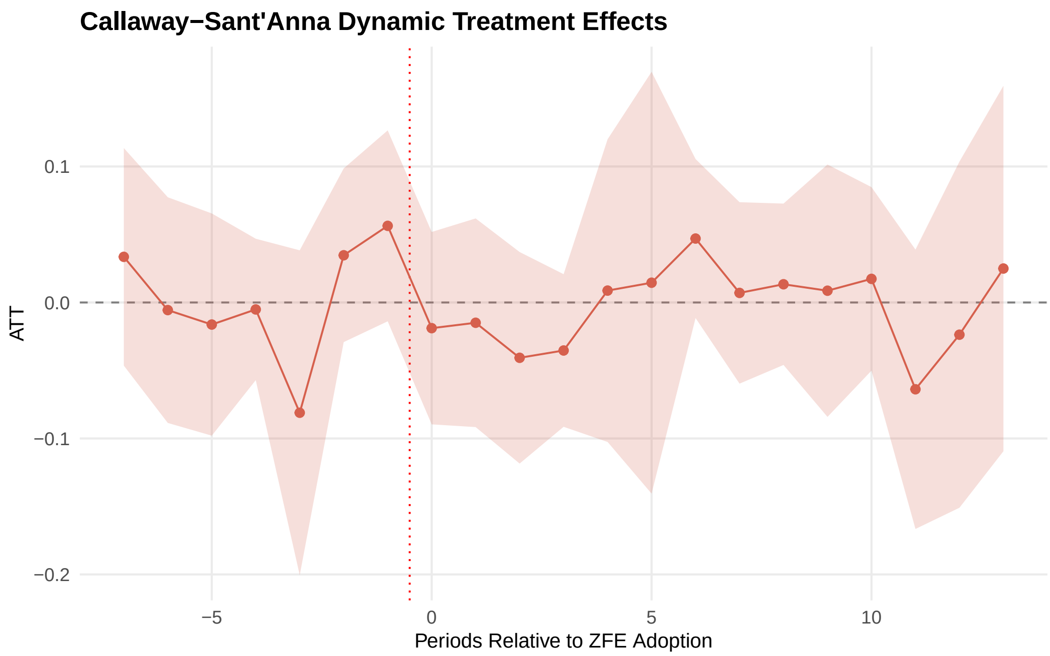

Figure 2: Callaway–Sant'Anna Dynamic Treatment Effects

Notes: Group-time average treatment effects aggregated by relative time, using Callaway and Sant'Anna (2021) with not-yet-treated comparison group. Shaded area shows 95% confidence intervals.

Figure 2 shows the dynamic Callaway–Sant'Anna effects. Unlike the TWFE event study, the pre-treatment effects are centered on zero, confirming that the CS-DiD successfully purges the confounding urban–suburban trends. Post-treatment effects are also indistinguishable from zero at all horizons.

5.4 Why TWFE Fails: The Urban–Suburban Confound

The divergence between TWFE (+11–22%) and CS-DiD (−0.3%) is explained by a simple fact: ZFE boundaries coincide with the urban–suburban divide.

In France, ZFEs cover city centers—dense, apartment-dominated neighborhoods with high demand. The "control" properties just outside the boundary are in suburban communes with different housing market dynamics. During 2020–2024, French city centers experienced faster price appreciation than suburbs, driven by post-COVID urban revaluation, public investment, and shifting preferences. The TWFE attributes this city-center premium to the ZFE, but the CS-DiD—by comparing across cities with different adoption timing—reveals that it is common to all metropolitan areas regardless of when (or whether) they adopted.

Three additional pieces of evidence support this interpretation:

Commercial diagnostic. A falsification test on commercial properties within the same boundary bandwidth produces a coefficient of 0.164 (SE = 0.055)—positive and significant. While commercial property values may respond to traffic restrictions and access changes, they should not respond to residential air quality amenity capitalization. The positive "effect" is more consistent with the TWFE capturing general urban trends than ZFE-specific residential amenity capitalization.

Distance gradient. Figure 3 shows treatment effects by distance ring. The gradient is monotonically increasing from outside to deep inside—but this is equally consistent with an urban price gradient as with amenity capitalization. The effect is positive even 2–5 km outside the boundary (relative to the reference), further suggesting a city-center premium rather than a boundary discontinuity.

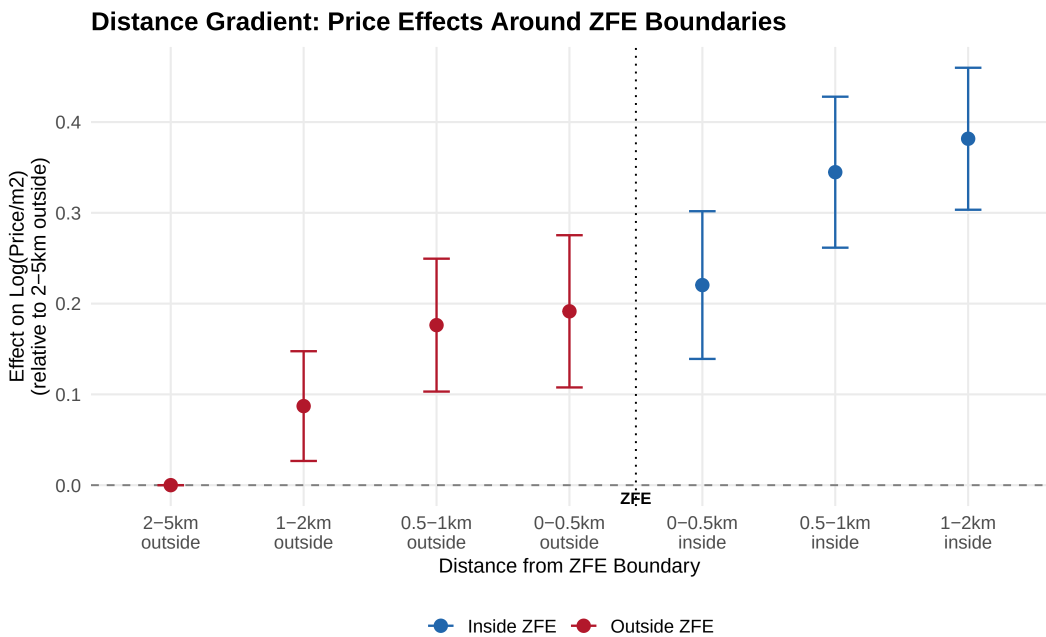

Figure 3: Distance Gradient: Price Effects Around ZFE Boundaries

Notes: Coefficients from a regression of log(price/m^2) on interactions between distance ring indicators and post-ZFE, with city and year-quarter fixed effects. Reference: 2–5 km outside. Error bars: 95% CI.

City-level heterogeneity. The TWFE effect varies enormously across cities (Figure 6): Paris shows +29%, Montpellier shows −21%, while Grenoble and Saint-Étienne are near zero. This pattern reflects each city's urban–suburban price dynamics rather than ZFE stringency or enforcement intensity.

Table 4: First Stage: ZFE Adoption and Air Quality

	(1) NO_2 (μg/m^3)	(2) $PM_{2.5}$ (μg/m^3)
Post ZFE	-0.42	-0.39*
	(0.51)	(0.20)
City FE	Yes	Yes
Year FE	Yes	Yes
Month-of-year FE	Yes	Yes
Observations	540	540

Notes: Unit of observation is city-month for the 9 cities in the boundary sample, 2020–2024. Dependent variable is mean monthly concentration. Data from CAMS reanalysis via Open-Meteo API. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

5.5 Air Quality First Stage

Table 4 presents the air quality first stage, estimated on a balanced panel of 9 cities $\times$ 60 months (2020–2024, $N = 540$). The policy barely moved the needle on air quality. NO_2—the primary target of vehicle bans—remained essentially flat (-0.42 μg/m^3, SE = 0.51, not significant). $PM_{2.5}$ showed a marginally significant reduction (-0.39 μg/m^3, SE = 0.20, $p = 0.056$). The specifications include city, year, and month-of-year fixed effects to absorb seasonal pollution cycles. The weak NO_2 first stage is consistent with imperfect enforcement and the gradual Crit'Air tightening that characterizes early ZFE implementation. If air quality barely improves, there is little amenity gain to capitalize into housing prices. Note that the city-centroid CAMS data (10 km resolution) measure city-level averages and cannot capture the within-city pollution gradient at the ZFE boundary; these results should be interpreted as suggestive rather than definitive.

6. Robustness and Heterogeneity

6.1 Bandwidth Sensitivity

Table 5 and Figure 7 show bandwidth sensitivity using the full geocoded sample (all transactions within each bandwidth, re-estimated from scratch). The TWFE coefficient increases with bandwidth—from 11 percent at 500 meters to 30 percent at 5 kilometers. The slight difference between the 2 km row in Table 5 and the baseline in Table 3 reflects minor sample differences in the geocoding match. In a clean boundary design, one would expect stability or slight increases at narrower bandwidths (where the comparison is tighter). The monotone increase with bandwidth is instead consistent with the TWFE capturing an urban–suburban

Table 5: Robustness Checks

Specification	Coefficient	SE	N
Bandwidth = 0.5km	0.1145***	(0.0287)	95,298
Bandwidth = 1km	0.1577***	(0.0295)	204,088
Bandwidth = 2km	0.2072***	(0.0304)	361,528
Bandwidth = 5km	0.3017***	(0.0333)	605,215
Donut (>200m from boundary)	0.2140***	(0.0311)	328,591
Placebo: Commercial properties	0.1642***	(0.0376)	27,732
Randomization inference p-value: 0.080 (500 permutations)			

Notes: All bandwidth specifications use city and year-quarter fixed effects (not commune FE), and hedonic controls (surface area, number of rooms). Coefficients may differ from Table 3 where commune FE are used. The dependent variable is log(price/m^2) for residential transactions. Standard errors clustered at the commune level. The placebo test uses commercial and industrial properties within 2km of ZFE boundaries. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

gradient that strengthens with distance from the center.

6.2 Donut Specification

Excluding transactions within 200 meters of the boundary produces a coefficient of 0.214 (SE = 0.031), nearly identical to the baseline TWFE. This rules out boundary sorting as the explanation for the positive TWFE but does not address the pre-trends concern.

6.3 CS-DiD Robustness: Leave-One-City-Out

To assess sensitivity to individual cities, I re-estimate the CS-DiD dropping one city at a time. The ATT ranges from −0.031 (dropping Paris) to +0.026 (dropping Rouen), with no estimate statistically distinguishable from zero. The null result is not driven by any single city—reassuring given the small number of treatment cohorts (seven) in the CS-DiD.

6.4 Randomization Inference

Permuting ZFE adoption dates across the nine boundary-sample cities 500 times yields a randomization inference p-value of 0.08 for the TWFE estimate. The mean permuted coefficient is 0.18—close to the actual estimate of 0.22—confirming that most of the TWFE "treatment effect" is present even under random adoption timing. This is direct evidence that the TWFE estimate reflects cross-city variation in urban–suburban trends rather than a causal ZFE effect.

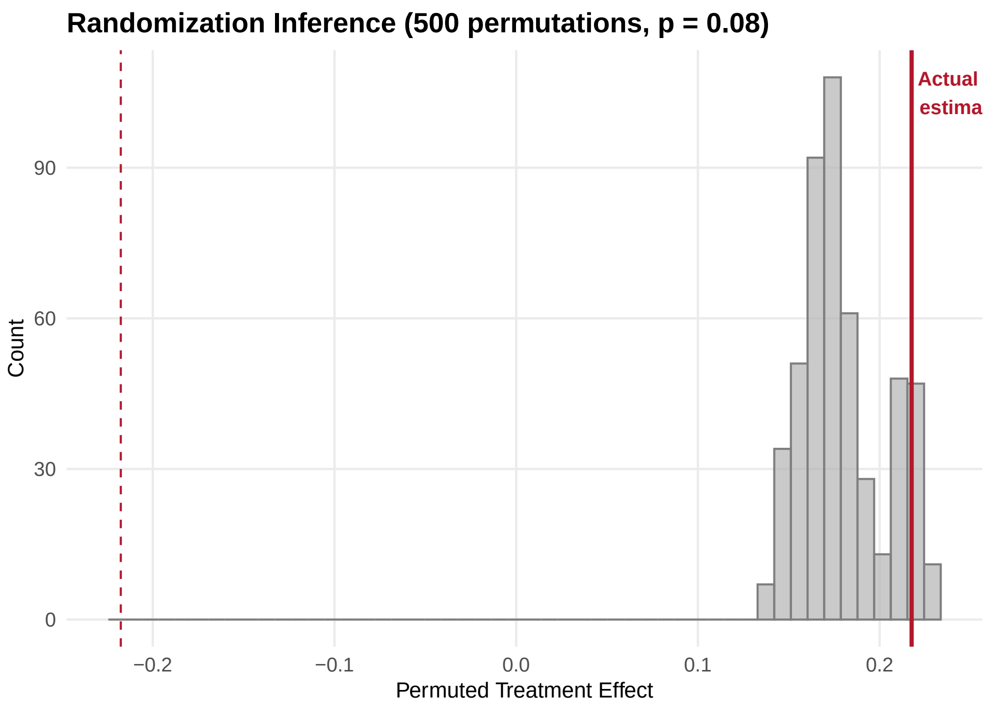

Figure 4: Randomization Inference Distribution

Notes: Distribution of treatment effects from 500 permutations of ZFE adoption dates. Red line: actual estimate. The RI p-value is 0.08. The high mean of the permuted distribution (0.18) indicates that most of the TWFE "effect" is non-causal.

6.5 Heterogeneity by Property Size

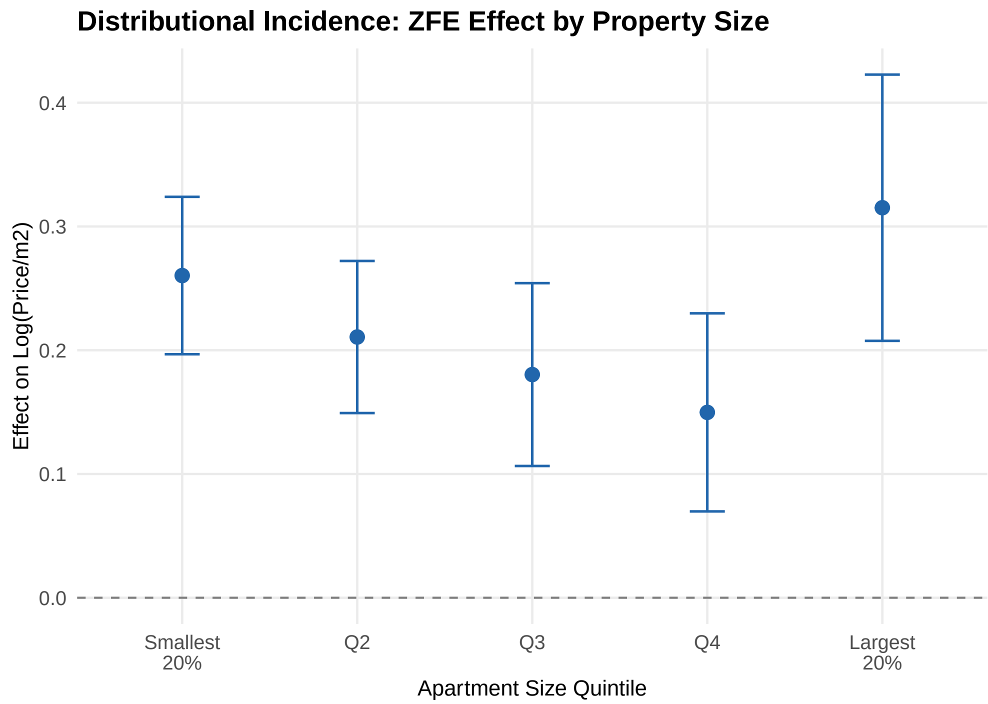

Figure 5: TWFE Effect by Apartment Size Quintile

Notes: Coefficients from TWFE regression with size quintile interactions. Given the pre-trends documented above, these should be interpreted as heterogeneous urban–suburban price dynamics, not causal ZFE effects.

Figure 5 shows the TWFE effect by apartment size quintile. The smallest apartments show the largest positive coefficient, with a negative interaction term for small apartments (-0.073, $p = 0.025$). However, given the pre-trends documented above, this heterogeneity should be interpreted with caution. It likely reflects that small urban apartments appreciated faster than small suburban apartments during this period—a pattern driven by post-COVID demand for compact urban living rather than ZFE-specific effects.

6.6 Heterogeneity by City

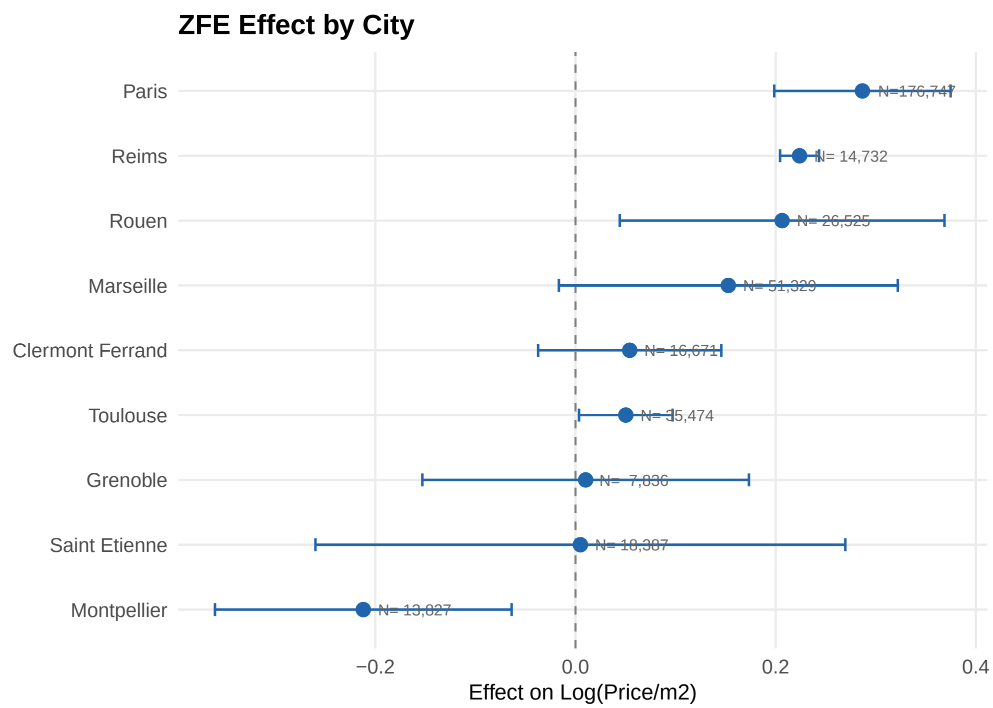

Figure 6: TWFE Effect by City

Notes: City-specific TWFE regressions. Error bars: 95% CI. The cross-city variation primarily reflects urban–suburban price dynamics, not ZFE policy variation.

Figure 6 reveals dramatic cross-city variation: Paris (+29%) and Reims (+22%) show large positive coefficients, while Montpellier (−21%) is strongly negative. This pattern bears no relationship to ZFE stringency, enforcement, or air quality improvements—Paris adopted its ZFE first and has the strictest enforcement, but so does Grenoble, which shows near-zero effect. The variation instead reflects city-specific urban–suburban price dynamics during 2020–2024.

7. Discussion

7.1 A Methodological Cautionary Tale

The central finding of this paper is a methodological one. The same data, same research design, and same bandwidth produce dramatically different estimates depending on the econometric approach: TWFE yields +22 percent; Callaway–Sant'Anna yields −0.3 percent. The event study reveals that the TWFE estimate is driven entirely by pre-existing differential trends.

This result has implications beyond ZFEs. Boundary difference-in-differences designs are popular in urban economics—used to study school quality (Black, 1999; Bayer et al., 2007), environmental amenities (Chay and Greenstone, 2005), and neighborhood effects (Banzhaf and Walsh, 2008). These designs are credible when the boundary is arbitrary with respect to the outcome variable. But many policy boundaries—low-emission zones, enterprise zones, school districts, opportunity zones—follow infrastructure or administrative borders that also delineate housing market segments. In such settings, the "boundary" confounds treatment with urban structure, and TWFE can produce large, significant, and wrong estimates.

The lesson is not to abandon boundary designs but to insist on rigorous pre-trend tests and estimators robust to staggered adoption. In this study, the Callaway–Sant'Anna estimator delivers a credible null by leveraging cross-city variation in adoption timing—precisely the variation that TWFE fails to exploit properly.

7.2 Why No Capitalization?

The near-zero estimate deserves substantive interpretation. Several explanations are consistent with the absence of detectable housing price effects:

Weak enforcement. ZFE fines are modest (68 euros) and enforcement is sporadic. If compliance is low and air quality improvements are small, there is little amenity gain to capitalize. The weak NO_2 first stage is consistent with this: ZFE adoption does not significantly reduce the primary vehicular pollutant.

Political uncertainty. Throughout 2022–2025, the future of ZFEs was politically uncertain. Property markets capitalize expected future amenities, and if buyers anticipated potential repeal, they would discount the expected ZFE benefit. The eventual repeal vote in January 2025 validates this concern.

Offsetting effects. ZFEs could simultaneously improve air quality (positive amenity) while restricting vehicle access (negative for car-dependent residents). If these effects offset, the net capitalization would be small. This is particularly plausible for the initial phase, when only the oldest vehicles were banned.

Measurement horizon. The analysis covers at most 3–5 years of post-adoption data for late-adopting cities. If capitalization is slow—requiring years for air quality improvements to accumulate and become salient—the null may reflect insufficient time rather than a true zero effect.

7.3 Policy Implications

The null result has direct policy relevance. The political debate over ZFEs centered on two distributional concerns: (i) the direct cost of vehicle compliance, which falls disproportionately on lower-income households, and (ii) the indirect cost of rising housing prices in ZFE neighborhoods, which could price out incumbent residents. This paper finds no evidence for the second concern. Whatever the merits of the vehicle compliance argument, the housing channel does not appear to be operative—at least not at the scale or precision that the data can detect.

This does not mean ZFEs have no distributional consequences. The vehicle compliance burden is real and documented. But the specific fear that ZFEs create large "green gentrification" effects through rising property values finds no support in the sale-price data for these later-adopting cities over the observed horizon. The evidence does not speak to rental markets, neighborhood composition changes, or longer-run effects.

7.4 Comparison with the Environmental Capitalization Literature

The null result stands in contrast to the positive capitalization effects documented in the U.S. environmental economics literature. Chay and Greenstone (2005) find that a one-unit reduction in TSP increases housing values by 0.2–1.5%, while Currie et al. (2015) document that toxic plant openings and closings affect property values within 1 mile by 1.5–3%. Greenstone et al. (2010) find large plant openings affect local property values by 1–2%.

Several differences may explain why ZFEs produce no detectable capitalization while U.S. environmental policies do. First, the environmental change is fundamentally different in scale. Superfund cleanups address severe contamination; ZFEs produce modest reductions in vehicular emissions. The $PM_{2.5}$ reduction I estimate ($-0.39\ \mu g/m^3$) is small relative to the baseline level—a roughly 5% reduction—which Ito and Zhang (2020) would predict generates only a modest willingness to pay.

Second, the policy permanence differs. U.S. environmental cleanups are largely irreversible; ZFEs were politically contested from inception and eventually repealed. Bajari et al. (2005) show that rational expectations about policy durability affect capitalization—if buyers discount the expected duration of a policy, the price effect is attenuated proportionally.

Third, information and salience differ. Superfund site listings generate media attention and are publicly disclosed; ZFE air quality improvements are diffuse and imperceptible to individual residents. The literature on the willingness to pay for clean air (Ito and Zhang, 2020; Deryugina et al., 2019) emphasizes that salience matters: people respond more to visible or dramatized pollution events than to statistical improvements in ambient concentrations.

7.5 External Validity

France's ZFE experience may not generalize to all low-emission zones. Several features of the French context—weak enforcement, modest fines, gradual phase-in, political uncertainty—could attenuate capitalization relative to more stringent regimes. London's ULEZ, for example, has stronger enforcement and higher fees (12.50 pounds per day), which may generate more salient effects. German low-emission zones, studied by Gehrsitz (2017) for health effects and Wolff (2014) for vehicle fleet effects, have been in place longer and may show different housing market dynamics.

The 2020–2024 analysis window also coincides with extraordinary macroeconomic conditions: the COVID-19 pandemic, a European energy crisis, and significant monetary policy shifts. These events may have dominated housing market dynamics, swamping any ZFE-specific signal. Future research using longer time series—particularly if some cities maintain their ZFEs despite the repeal mandate—would provide more definitive evidence.

7.6 Limitations

Several limitations merit acknowledgment and should inform the interpretation of results.

Limited pre-treatment period. Geocoded DVF data are available only from 2020. Paris adopted its ZFE in 2017 and Grenoble in 2019, meaning these cities are already treated when the data begin and are excluded from the Callaway–Sant'Anna estimation. The CS-DiD therefore identifies effects from the seven cities adopting during 2020–2023 (Rouen through Reims), comparing them to not-yet-treated cities at each point in time. This restriction limits the number of treated cohorts and may miss effects specific to the earliest adopters. The limited pre-treatment horizon also restricts statistical power for pre-trend tests and may miss slow-moving anticipation effects.

Coarse air quality data. The CAMS reanalysis data have a resolution of approximately 10 km—sufficient for comparing city-level trends but too coarse to capture within-city pollution gradients. If air quality improves more inside the ZFE core than near the boundary (which is likely given traffic patterns), the city-level first stage understates the true pollution reduction experienced by treated properties. With higher-resolution monitoring data (e.g., from ATMO France's ground stations), one could test whether within-city pollution gradients correlate with property price changes.

Aggregation for CS-DiD. The Callaway–Sant'Anna estimator requires a balanced panel, which necessitates aggregating transaction-level data to commune-quarter cells. This aggregation discards within-commune, within-quarter variation and reduces the sample from hundreds of thousands of transactions to a few thousand commune-quarter observations. The resulting loss of precision could contribute to the null result. However, the CS-DiD confidence interval ($\pm 5\%$) is tight enough to rule out economically large effects.

No rental data. This study uses transaction prices, not rents. Rental markets may respond differently to amenity changes, particularly if renters are more mobile than owners. If ZFEs increase rents without affecting transaction prices (e.g., because of rent regulation interactions), the welfare effects could be substantial even with zero price capitalization. France's rental price indices (Clameur, SeLoger) are proprietary and not available for this analysis, but future work integrating rental data would complement these findings.

Incomplete ZFE coverage. Strasbourg, one of the larger French metropolitan areas with an active ZFE, is excluded because DVF does not cover the Alsace-Moselle departments. Additionally, Nice and Lyon adopted ZFEs that initially covered limited areas, reducing the number of treated transactions for these cities. The results are driven primarily by Paris (which accounts for 49% of the boundary sample), limiting the generalizability of city-level heterogeneity estimates.

8. Conclusion

France's low-emission zones were adopted amid promises of cleaner air and fears of rising housing costs. Using the universe of geocoded property transactions and official ZFE boundary polygons, this paper finds no evidence of large capitalization effects—at least through the sale-price channel in later-adopting cities. A robust staggered difference-in-differences estimator yields a near-zero effect of −0.3 percent (95% CI: −5.2 to +4.6 percent), ruling out large effects while remaining consistent with small positive or negative impacts.

The paper also delivers a methodological warning. Naive TWFE estimates suggest a 22 percent price premium, significant at conventional levels and large enough to dominate any policy discussion. But this estimate is an artifact of ZFE boundaries coinciding with the urban–suburban divide. Pre-trend tests, commercial diagnostics, and randomization inference all point to the same conclusion: the TWFE captures city-center appreciation, not ZFE capitalization.

For policymakers, the null result offers a silver lining. The most politically potent argument against low-emission zones—that they substantially raise property values, pricing out incumbent residents—finds no support in the sale-price data for France's later-adopting cities during 2020–2024. The confidence interval rules out effects larger than roughly 5 percent. But the null may also reflect a policy that, in practice, changed little: modest fines, sporadic enforcement, and gradual phase-ins may have been insufficient to produce measurable effects on either air quality or property values. The evidence does not speak to rental markets or longer-run dynamics.

The broader lesson extends to environmental policy design. Place-based environmental improvements *can* capitalize into property values, creating winners (owners) and losers (renters). But capitalization requires that the policy produce a perceptible and durable amenity improvement. When implementation is tentative, enforcement is weak, and political commitment is uncertain, the hedonic channel may stay dormant—which is good news for equity, even if it raises questions about policy effectiveness.

The methodological lesson is equally important. Boundary difference-in-differences designs have become a workhorse of urban and environmental economics since Black (1999). This paper shows that these designs can fail dramatically when the boundary of interest coincides with systematic differences in housing market dynamics. The TWFE estimate of +22 percent would dominate any policy discussion and was precisely estimated with conventional standard errors. Yet it was entirely driven by pre-existing urban–suburban trends. The Callaway–Sant'Anna estimator, by exploiting variation in adoption timing across cities, recovers the credible null that the TWFE obscures.

Researchers applying boundary designs to policy zones—low-emission zones, opportunity zones, enterprise zones, school districts—should routinely report both TWFE and staggered estimators, present detailed pre-trend diagnostics, and conduct diagnostic tests on outcomes less likely to respond to the specific policy channel. The divergence documented here is not an exotic pathology; it is a predictable consequence of policy boundaries following urban infrastructure.

Acknowledgements

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

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

Contributors: @olafdrw

References

Bajari, Patrick, C. Lanier Benkard, and John Krainer, "A rational expectations approach to hedonic price regressions with time-varying unobserved product attributes," *American Economic Review*, 2005, *95* (1), 258–289.

Banzhaf, H. Spencer and Randall P. Walsh, "Do people vote with their feet? An empirical test of Tiebout," *American Economic Review*, 2008, *98* (3), 843–863.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan, "A unified framework for measuring preferences for schools and neighborhoods," *Journal of Political Economy*, 2007, *115* (4), 588–638.

Bergé, Laurent, "Efficient estimation of multi-way fixed effects models," *Journal of Econometrics*, 2018. R package `fixest`.

Black, Sandra E., "Do better schools matter? Parental valuation of elementary education," *Quarterly Journal of Economics*, 1999, *114* (2), 577–599.

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

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

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-based improvements for inference with clustered errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Chay, Kenneth Y. and Michael Greenstone, "Does air quality matter? Evidence from the housing market," *Journal of Political Economy*, 2005, *113* (2), 376–424.

Currie, Janet, Lucas Davis, Michael Greenstone, and Reed Walker, "Environmental health risks and housing values: Evidence from 1,600 toxic plant openings and closings," *American Economic Review*, 2015, *105* (2), 678–709.

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

Deryugina, Tatyana, Garth Heutel, Nolan H. Miller, David Molitor, and Julian Reif, "The mortality and medical costs of air pollution: Evidence from changes in wind direction," *American Economic Review*, 2019, *109* (12), 4178–4219.

Gehrsitz, Markus, "The effect of low-emission zones on air pollution and infant health," *Journal of Environmental Economics and Management*, 2017, *83*, 121–144.

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

Green, Colin P., John S. Heywood, and Maria Navarro Paniagua, "Low-emission zones and road pricing: Clean Air Zones in the UK," *Journal of Environmental Economics and Management*, 2020, *100*, 102287.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti, "Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings," *Journal of Political Economy*, 2010, *118* (3), 536–598.

Holman, Claire, Roy Harrison, and Xavier Querol, "Where do Low Emission Zones help? Evidence from London," *Atmospheric Environment*, 2015, *112*, 193–201.

Ito, Koichiro and Shuang Zhang, "The willingness to pay for clean air: Evidence from air purifier markets in China," *Journal of Political Economy*, 2020, *128* (5), 1627–1672.

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

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

Wolff, Hendrik, "Keep your clunker in the suburb: Low-emission zones and adoption of green vehicles," *The Economic Journal*, 2014, *124* (578), F481–F512.

A. Data Appendix

A.1 DVF Data Processing

The DVF (*Demandes de Valeurs Foncières*) data are downloaded from files.data.gouv.fr/geo-dvf in geocoded CSV format, organized by department and year. Each file contains the universe of property transactions in that department-year. Geocoding was performed by Etalab using the Base Adresse Nationale (BAN) and cadastral parcel centroids.

Sample restrictions applied sequentially:

1. **Transaction type:** Retain only *Vente* (arm's-length sales).
2. **Price:** Drop transactions below 10,000 or above 10,000,000 euros.
3. **Surface:** Drop below 9 m^2 or above 1,000 m^2.
4. **Price per m^2:** Drop below 500 or above 30,000 euros.
5. **Geocoding:** Drop transactions without valid coordinates.

A.2 ZFE Boundary Processing

The BNZFE is maintained by the DGITM and published on transport.data.gouv.fr. Each ZFE phase includes:

- `date_debut`: Phase effective date
- `vp_critair`: Crit'Air threshold for passenger vehicles
- `vul_critair`: Threshold for light commercial vehicles
- `pl_critair`: Threshold for heavy vehicles

The multiple BNZFE features per city represent successive Crit'Air threshold tightenings within a geographically stable perimeter, not geographic expansions of the zone. All ZFE polygons are unioned into a single multi-polygon for spatial classification. Each transaction is classified as inside/outside via point-in-polygon, and signed distance to the boundary is computed in meters using the Lambert-93 projection (EPSG:2154). Treatment timing is assigned using each city's first ZFE adoption date (Table 1), ensuring that a property is coded as treated only after its city activated a ZFE.

B. Identification Appendix

B.1 Callaway–Sant'Anna Implementation

I aggregate transaction-level data to commune-quarter cells (mean log price/m^2) before estimating the Callaway and Sant'Anna (2021) model. The treatment assignment mirrors Section 4: communes inside ZFE boundaries receive the group indicator G_i equal to their city's adoption quarter; communes outside ZFE boundaries receive $G_i = 0$ (never-treated). The control group at each period consists of all never-treated communes (outside-boundary) plus inside-boundary communes in not-yet-treated cities. Paris and Grenoble are excluded because their inside-boundary communes are already treated when the 2020–2024 sample begins and thus lack pre-treatment observations. The varying base period option allows group-specific reference periods.

Group-time ATTs are aggregated in two ways: (i) simple overall ATT averaging across all (g, t) cells, and (ii) dynamic aggregation by relative period.

B.2 Randomization Inference

I conduct randomization inference by permuting ZFE adoption dates across the nine boundary-sample cities 500 times. For each permutation, I randomly reassign these nine cities' adoption dates and re-estimate the TWFE hedonic specification. The RI p-value is the fraction of permuted coefficients with absolute value exceeding the actual estimate. The high mean of the permuted distribution (0.180 vs. actual 0.218) confirms that most of the TWFE signal is non-causal.

C. Additional Figures

C.1 Bandwidth Robustness

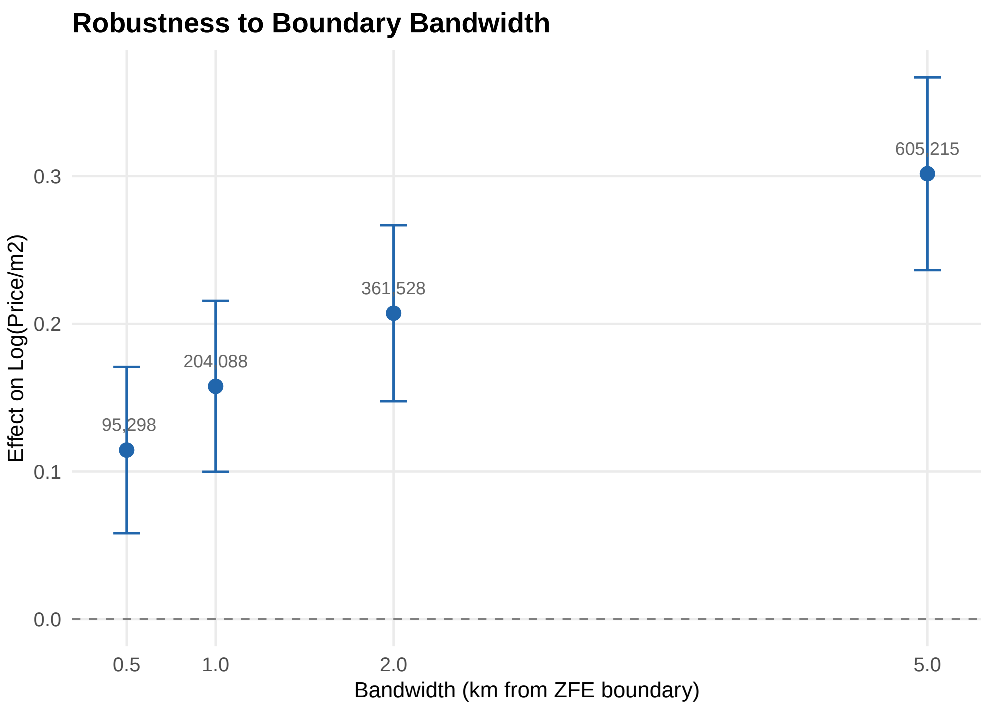

Figure 7: Robustness to Boundary Bandwidth

Notes: TWFE coefficient by bandwidth. Numbers above points indicate sample size. The monotone increase with bandwidth is consistent with an urban–suburban gradient.

C.2 Air Quality Trends

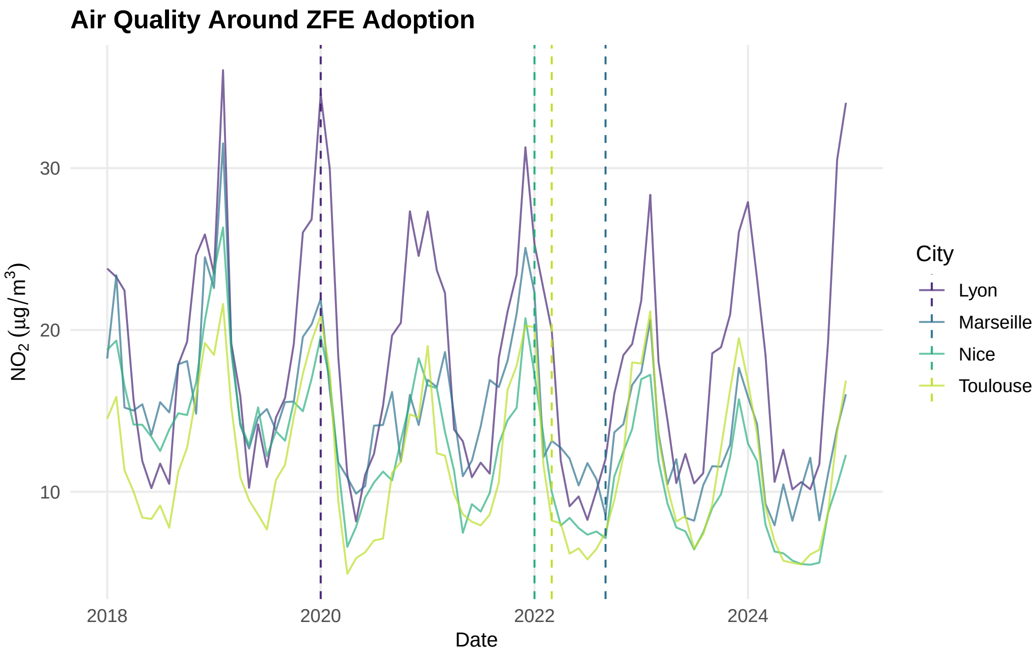

Figure 8: NO_2 Trends Around ZFE Adoption

Notes: Monthly mean NO_2 for selected cities. Vertical dashed lines: ZFE adoption dates. Data: CAMS reanalysis via Open-Meteo.