

Clean Air, Dirty Divide? Property Price Effects of Low Emission Zone Boundaries in France

APEP Autonomous Research* @olafdrw

February 25, 2026

Abstract

Low emission zones restrict polluting vehicles within designated urban perimeters, but their effects on property markets remain unstudied. I exploit the sharp geographic boundary of France's Grand Paris *Zone à Faibles Émissions* (ZFE) using a spatial regression discontinuity design, comparing geocoded property transactions just inside versus outside the A86 motorway boundary during the strong enforcement period (June 2021 onward, Crit'Air 4+). I find no statistically significant price discontinuity at the ZFE boundary (-2.4% , $p = 0.45$). The null persists across bandwidths, polynomial orders, and enforcement periods. These results suggest that either ZFE environmental benefits are too diffuse to capitalize at the boundary, enforcement is too weak to alter perceived neighborhood quality, or offsetting forces—accessibility losses versus air quality gains—neutralize the net price effect.

JEL Codes: Q53, R31, R48, H23

Keywords: low emission zones, property values, spatial regression discontinuity, air quality, housing markets, France

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

1. Introduction

In Paris, the A86 motorway marks a silent border. On one side, an old diesel Peugeot faces a €68 fine; on the other, it roams freely. As over 300 cities worldwide draw these invisible lines to banish pollution (Gehrsitz, 2017), they are also rewriting the map of urban wealth—or so the theory goes. Do these borders actually reshape housing markets?

The question matters for two reasons. First, if property values respond to low emission zone boundaries, it reveals how households value the environmental amenity these zones create. The hedonic pricing framework, pioneered by Rosen (1974) and applied to environmental goods by Chay and Greenstone (2005) and Currie et al. (2015), teaches us that housing prices capitalize local amenities. A price premium inside a low emission zone would indicate that buyers perceive meaningful air quality improvements worth paying for. Second, and more troublingly, any price effect at the boundary implies that low emission zones create a new dimension of spatial inequality. If cleaner air inside the zone raises property values, it may price out the very households who depend most on car access—lower-income residents with older, more polluting vehicles.

This paper provides the first within-city boundary evidence on how low emission zones affect property values, and the first evidence from France. I study France’s *Zones à Faibles Émissions* (ZFE), which have been progressively implemented across major cities since 2019 under the *Loi d’Orientation des Mobilités* and expanded by the *Loi Climat et Résilience* of 2021. I focus on the Metropole du Grand Paris ZFE, whose boundary follows the A86 motorway—a pre-existing road that was not drawn for housing market reasons, providing a clean spatial discontinuity.

My identification strategy is a spatial regression discontinuity design (spatial RDD). The running variable is the signed geodesic distance from each property transaction to the nearest point on the ZFE boundary, measured in kilometers. Properties with positive distance lie inside the zone; those with negative distance lie outside. I estimate local polynomial regressions at the boundary following Calonico et al. (2014) and Cattaneo et al. (2020b), using MSE-optimal bandwidth selection and robust bias-corrected inference.

The data come from two public sources. France’s *Demandes de Valeurs Foncières* (DVF) provides the universe of property transactions with geocoded coordinates, covering over three million sales per year with 98.6% geocoding accuracy. I combine this with GeoJSON boundary polygons for the Paris ZFE, downloaded from the city’s open data portal. The combination allows me to compute precise distances from each transaction to the zone boundary and construct a dataset uniquely suited to spatial RDD estimation.

The ZFE boundary does not move property prices. The point estimate is -0.024 log points

(-2.4% , $p = 0.45$)—economically small and statistically indistinguishable from zero. This null persists across specifications with property controls, alternative bandwidths, polynomial orders, and donut hole designs. A weak-enforcement placebo using 2020 transactions also shows no discontinuity.

Several features complicate interpretation. The McCrary density test (McCrary, 2008; Cattaneo et al., 2020a) rejects smooth density at the boundary, reflecting the A86 motorway’s role as a pre-existing barrier that shapes urban development asymmetrically. The housing stock differs across the boundary: apartments are more prevalent inside (denser urban fabric) while houses predominate outside. While continuous covariates like surface area and rooms are balanced, the compositional shift in building type suggests the A86 creates a genuine urban discontinuity that is not solely attributable to the ZFE.

The heterogeneity analysis reveals interesting patterns. Apartments show a precisely estimated null (-0.020 , $p = 0.53$), while houses show a large but imprecise positive estimate (0.46 , $p = 0.11$). Year-by-year estimates show no clear temporal pattern: a positive effect appears in 2021 ($+0.080$, $p = 0.01$), but subsequent years revert to null. This lack of a monotonic enforcement-response pattern is inconsistent with simple capitalization of progressively tighter restrictions.

This paper contributes to three literatures. First, it advances the environmental valuation literature that uses hedonic methods to estimate willingness to pay for clean air (Chay and Greenstone, 2005; Currie et al., 2015). While previous work has exploited plant openings, Superfund cleanups, and pollution monitors as sources of variation, I am the first to use a regulatory boundary that directly separates polluted and less-polluted areas by design. The spatial RDD provides a sharper identification strategy than the difference-in-differences approaches used in prior low emission zone studies (Pestel and Wozny, 2025; Aydin and Kürschner Rauck, 2023).

Second, I contribute to the growing literature on low emission zones, which has focused almost exclusively on their effects on air quality (Gehrsitz, 2017), vehicle fleet composition (Wolff, 2014), and, recently, housing rents in Germany (Pestel and Wozny, 2025). All existing property market studies use difference-in-differences designs comparing cities with and without zones. By exploiting the spatial boundary *within* a city, I eliminate city-level confounders and provide the first evidence from France, where the ZFE policy is among Europe’s most ambitious.

Third, I speak to the broader question of how place-based environmental policies create or exacerbate spatial inequality (Tiebout, 1956; Banzhaf and Walsh, 2008). If low emission zones raise property values inside the boundary, they effectively create a “green premium” that benefits existing homeowners but raises barriers to entry for renters and prospective

buyers. The distributional implications are particularly salient in France, where political backlash against ZFE restrictions has been intense, echoing the *gilets jaunes* movement.

2. Institutional Background

2.1 France’s Low Emission Zones

France’s *Zones à Faibles Émissions mobilité* (ZFE-m) are geographically defined urban areas where vehicles failing to meet specified emission standards are prohibited. The legal framework builds on three pillars: the 2015 Energy Transition Law (*Loi de Transition Énergétique*), which created the Crit’Air vehicle classification system; the 2019 Mobility Orientation Law (*Loi d’Orientation des Mobilités*, LOM), which gave local authorities the power to create ZFEs; and the 2021 Climate and Resilience Law (*Loi Climat et Résilience*), which mandated ZFEs in all agglomerations exceeding 150,000 inhabitants where air quality standards are regularly exceeded.

The Crit’Air system classifies vehicles on a 0–5 scale based on engine type and registration date. Crit’Air 0 denotes zero-emission vehicles (electric, hydrogen). Crit’Air 1 covers recent petrol and hybrid vehicles. Higher numbers indicate older, more polluting vehicles, with Crit’Air 5 corresponding to diesel vehicles registered before 2001 and petrol vehicles before 1997. Vehicles without any Crit’Air sticker (*non classés*) are typically pre-1997 and face the strictest restrictions.

ZFE restrictions follow a progressive calendar. Local authorities set the pace, but the 2021 law established minimum milestones: by 2023, all mandatory ZFEs must ban at least Crit’Air 5 vehicles; by 2024, Crit’Air 4; and by 2025, Crit’Air 3. This progressive tightening means that the effective treatment intensity increases over time, creating useful temporal variation for studying how property values respond to enforcement stringency.

2.2 The Grand Paris Métropole ZFE

The Metropole du Grand Paris (MGP) ZFE is France’s largest, covering 79 communes and approximately 5.6 million residents. Its geographic boundary follows the A86 motorway, a 79-kilometer orbital road that encircles the city of Paris and its immediate suburbs (*petite couronne*: Hauts-de-Seine, Seine-Saint-Denis, Val-de-Marne). The A86 itself is excluded from the ZFE perimeter, functioning as the border.

The MGP ZFE has operated since June 1, 2019, initially banning Crit’Air 5 and unclassified vehicles. The restriction timeline is:

- **June 2019:** Crit’Air 5 and unclassified vehicles banned (diesel pre-2001, petrol pre-1997)
- **June 2021:** Crit’Air 4 vehicles banned (diesel pre-2006)
- **July 2023:** Crit’Air 3 vehicles banned (diesel pre-2011, petrol pre-2006)

Enforcement relies on police checks and, increasingly, automated camera surveillance (*contrôle automatisé*), though enforcement intensity has been debated. Penalties for violations are €68 for passenger cars and €135 for heavy vehicles.

The A86 boundary is particularly attractive for a spatial RDD because it was not drawn to reflect housing market characteristics. The motorway was constructed between 1969 and 2011, decades before any ZFE policy was contemplated. Its path was determined by transportation planning considerations—connecting ring suburbs and reducing radial traffic loads—rather than by residential property values or neighborhood amenities. This “pre-existing boundary” property strengthens the identifying assumption that the ZFE boundary provides as-good-as-random variation in treatment assignment for properties near the border.

3. Data

3.1 Property Transactions: DVF

I use France’s *Demandes de Valeurs Foncières* (DVF), a comprehensive registry of property transactions maintained by the Direction Générale des Finances Publiques (DGFIP) and published as open data. DVF records every property sale in metropolitan France with details on price, property characteristics, and location. Since 2019, the data have been geocoded by Etalab with latitude/longitude coordinates, achieving a 98.6% geocoding rate.

Each observation in DVF represents a transaction (*mutation*) and includes: the sale date (*date_mutation*), total price (*valeur_foncière*), building type (*type_local*: apartment or house), built surface area (*surface_reelle_bati*), number of principal rooms (*nombre_pieces_principales*), number of lots (*nombre_lots*), land area (*surface_terrain*), and geographic identifiers (commune code, department code, postal code).

I restrict the sample to residential sales of apartments and houses with non-missing prices, coordinates, and surface area. I compute price per square meter as the primary outcome variable and exclude extreme outliers below €500/m² or above €30,000/m², which are likely data errors or highly atypical commercial transactions.

The analysis sample covers departments 75 (Paris), 92 (Hauts-de-Seine), 93 (Seine-Saint-Denis), and 94 (Val-de-Marne)—the four departments that straddle the A86 boundary of the

Grand Paris ZFE. The sample period spans July 2020 through June 2025 using data from the publicly available DVF files on data.gouv.fr.¹ This provides post-ZFE observations across progressive enforcement phases from Crit’Air 5 (June 2019) through Crit’Air 3 (July 2023 onward).

3.2 ZFE Boundary Data

The geographic perimeter of the Grand Paris ZFE is obtained from the Paris Open Data portal as a GeoJSON polygon file. This polygon delineates the exact boundary of the ZFE, including all streets and parcels inside and outside the restricted area. I validate the polygon against official decree maps published in the *Journal Officiel*.

3.3 Spatial Distance Computation

For each DVF transaction, I compute the signed geodesic distance to the nearest point on the ZFE boundary polygon. I project all coordinates to Lambert-93 (EPSG:2154), France’s official projected coordinate reference system, to ensure accurate metric distances. The signed distance is positive for properties inside the ZFE and negative for properties outside, with zero at the boundary.

This signed distance serves as the running variable for the spatial RDD. The key assumption is that properties just inside and just outside the boundary are comparable in all respects except ZFE treatment status—an assumption I test formally in Section 5.

¹DVF data were downloaded from <https://files.data.gouv.fr/geo-dvf/latest/csv/> (departmental annual files for 2020–2023) and <https://files.data.gouv.fr/geo-dvf/latest/csv/full.csv.gz> (consolidated file including 2024–2025 transactions). Data accessed February 2026.

3.4 Summary Statistics

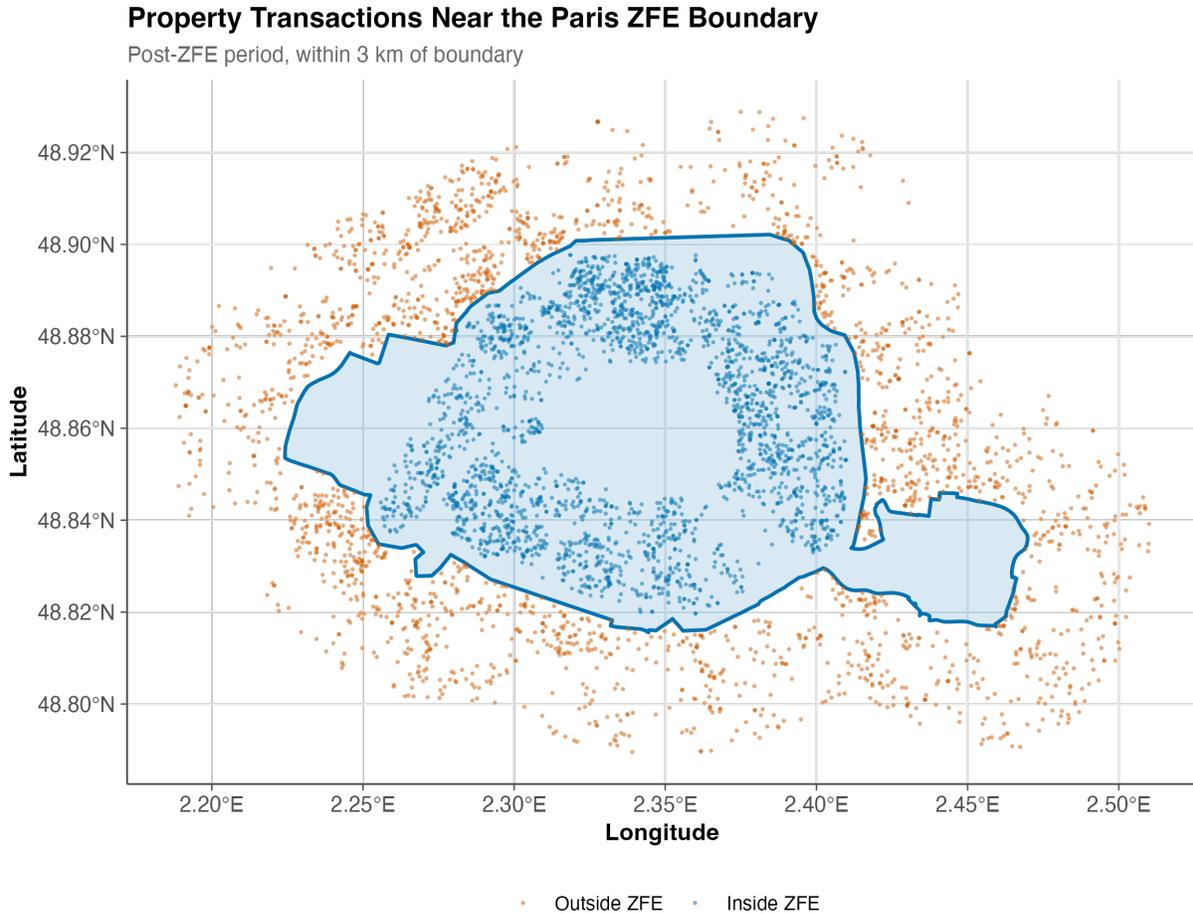


Figure 1: Property Transactions Near the Paris ZFE Boundary
Notes: Geocoded DVF transactions within 3 km of the Paris ZFE boundary. Blue: inside ZFE. Orange: outside ZFE. The shaded polygon shows the ZFE perimeter following the A86 motorway. Source: DVF (data.gouv.fr), Paris Open Data (ZFE boundary).

Figure 1 displays the geographic distribution of property transactions near the ZFE boundary. The A86 motorway's ring structure is clearly visible, with denser transaction clusters inside (the *petite couronne*) and more dispersed patterns outside.

Table 1: Summary Statistics

	Weak Enforcement		Strong Enforcement	
	Outside	Inside	Outside	Inside
Price/m ² (€)	6,462 (4,031)	11,442 (4,011)	6,514 (3,896)	10,991 (4,026)
Surface (m ²)	66.3	52.5	64.8	53.5
Rooms	3.06	2.4	2.99	2.41
Apartment (%)	77.3	99.4	78.6	99.5
Observations	29,158	15,455	202,783	127,713

Notes: Standard deviations in parentheses for price. “Inside” and “Outside” refer to the ZFE boundary (A86 motorway). “Weak Enforcement” = 2020 (Crit’Air 5 only). “Strong Enforcement” = June 2021 onward (Crit’Air 4+). Sample restricted to residential sales (apartments and houses) with price per m² between €500 and €30,000. Source: DVF (data.gouv.fr).

Table 1 presents summary statistics for the analysis sample, stratified by location (inside vs. outside the ZFE) and enforcement phase (weak: 2020, Crit’Air 5 only; strong: June 2021 onward, Crit’Air 4+). The main RDD estimation uses the strong enforcement subsample; the weak enforcement sample serves as a placebo. Properties inside the ZFE are, on average, smaller and more likely to be apartments, reflecting the denser urban fabric of inner Paris and the *petite couronne*. Prices per square meter are higher inside the ZFE in both phases, but this level difference reflects the well-known Paris price gradient and is not itself evidence of a ZFE effect. The RDD identifies the *discontinuity* at the boundary, not the level difference.

4. Empirical Strategy

4.1 Spatial Regression Discontinuity Design

I estimate the effect of the ZFE on property prices using a spatial regression discontinuity design. The identifying assumption is that the conditional expectation of potential outcomes under no treatment is continuous at the ZFE boundary:

$$\lim_{d \downarrow 0} \mathbb{E}[Y_i(0) | D_i = d] = \lim_{d \uparrow 0} \mathbb{E}[Y_i(0) | D_i = d] \quad (1)$$

where $Y_i(0) = \log(p_i/m_i^2)$ is the potential log price per square meter absent ZFE treatment and D_i is the signed distance to the boundary. Under this assumption, the treatment effect at the boundary is identified as:

$$\tau = \lim_{d \downarrow 0} \mathbb{E}[Y_i | D_i = d] - \lim_{d \uparrow 0} \mathbb{E}[Y_i | D_i = d] \quad (2)$$

The primary specification estimates local polynomial regressions of the form:

$$\log(p_i/m_i^2) = \alpha + \tau \cdot \mathbb{I}[D_i \geq 0] + f(D_i) + X_i' \beta + \gamma_t + \varepsilon_i \quad (3)$$

where $\mathbb{I}[D_i \geq 0]$ indicates that property i lies inside the ZFE, $f(D_i)$ is a local polynomial in signed distance (linear in the baseline, quadratic as robustness), X_i is a vector of property controls (surface area, rooms, apartment indicator), γ_t captures year-quarter fixed effects to account for seasonal and macroeconomic trends, and ε_i is the error term. The coefficient τ measures the log-point price discontinuity at the ZFE boundary.

I estimate this using the `rdrobust` package (Calonico et al., 2014, 2020), which implements MSE-optimal bandwidth selection with robust bias-corrected confidence intervals. The baseline specification uses a triangular kernel, which assigns greater weight to observations closer to the boundary. I report robust standard errors throughout.

4.2 Identification Assumptions and Threats

The validity of the spatial RDD rests on three assumptions, each of which I test directly.

No manipulation. Properties cannot strategically relocate across the boundary, so the standard manipulation concern in RDD—that agents sort around the threshold—is less relevant here. However, if development patterns differ systematically at the boundary, or if more transactions occur on one side, the density of the running variable could be discontinuous. I conduct a formal density test using the local polynomial method of Cattaneo et al. (2020a).

Continuity of potential outcomes. If the boundary follows a pre-existing feature that independently affects property values, the continuity assumption fails. The A86 motorway is such a feature: proximity to a major highway could depress property values through noise and pollution, creating a dip at the boundary unrelated to the ZFE. I address this in three ways. First, I test whether continuous property characteristics (surface area, rooms) are smooth at the boundary. Second, I acknowledge that compositional characteristics (apartment share, lot count) *do* change discontinuously, reflecting the pre-existing urban–suburban transition at the A86, and control for property type in the augmented specification. Third, I exploit temporal variation: I compare the boundary discontinuity during weak versus strong enforcement. If

the A86 itself depresses prices, this effect should be constant over time. A *change* in the boundary discontinuity with stronger enforcement can only be attributed to the ZFE.

No spillovers. If the ZFE displaces polluting vehicles to just outside the boundary, concentrating pollution there, this would amplify the estimated effect. Conversely, if improved air quality spills beyond the boundary, the estimate would be attenuated. I address this by estimating donut-hole specifications that exclude observations very close to the boundary.

4.3 Difference-in-Discontinuities

To isolate the ZFE effect from any pre-existing boundary discontinuity, I implement a difference-in-discontinuities estimator. Let $\hat{\tau}_{\text{strong}}$ denote the boundary discontinuity during strong enforcement (June 2021 onward, Crit’Air 4+) and $\hat{\tau}_{\text{weak}}$ the discontinuity during weak enforcement (2020, Crit’Air 5 only). The difference-in-discontinuities estimate is:

$$\hat{\tau}_{\text{DiD}} = \hat{\tau}_{\text{strong}} - \hat{\tau}_{\text{weak}} \tag{4}$$

This removes any time-invariant boundary effect (such as highway noise) and isolates the component attributable to the intensification of ZFE restrictions. Under the assumption that the A86’s effect on property prices did not change coincidentally with the enforcement tightening, $\hat{\tau}_{\text{DiD}}$ identifies the causal effect of stronger ZFE enforcement.

5. Results

5.1 Main Results

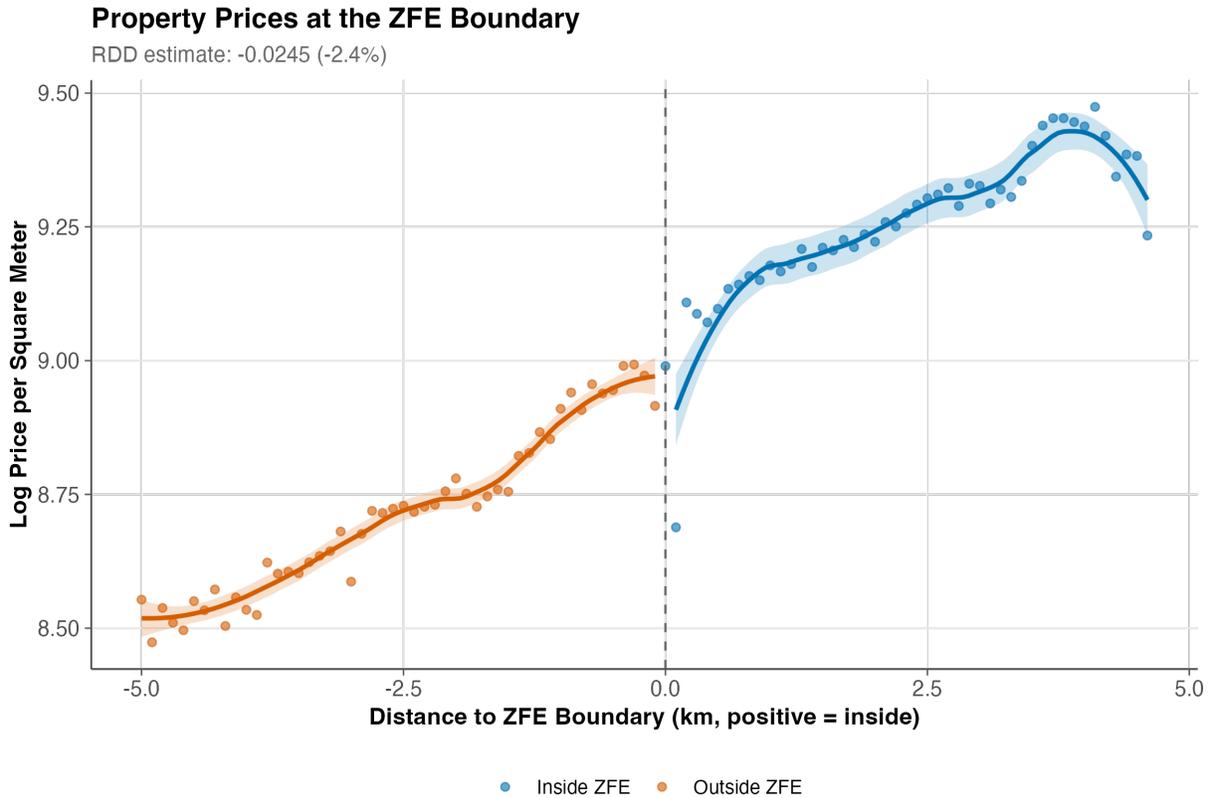


Figure 2: Property Prices at the ZFE Boundary

Notes: Each point represents the mean log price per m² within a 100-meter distance bin. Solid lines are local polynomial fits estimated separately on each side of the boundary with 95% confidence intervals. The dashed vertical line marks the ZFE boundary (distance = 0). Positive distances indicate properties inside the ZFE.

Figure 2 plots mean log prices per square meter in 100-meter bins against distance to the ZFE boundary, with local polynomial fits on each side. While raw price levels are higher inside the ZFE—reflecting the well-known Paris price gradient—no sharp discontinuity is visible at the boundary itself. The local polynomial fits on each side approach the cutoff at similar levels, foreshadowing the formal null result.

Table 2: Main Results: Effect of ZFE on Property Prices

	(1)	(2)	(3)
	Strong Enf.	Strong Enf.	Weak Enf.
	Baseline	+ Controls	Placebo
Inside ZFE	-0.0245 (0.0327)	-0.0184 (0.0328)	-0.0392 (0.0531)
Implied % effect	-2.42%	-1.82%	-3.85%
Controls	No	Yes	No
Bandwidth (km)	0.414	0.421	0.779
N (left)	202,783	202,783	29,158
N (right)	127,713	127,713	15,455

Notes: Robust bias-corrected standard errors in parentheses (Cattaneo, Idrobo, and Titiunik 2019). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Local linear regression with triangular kernel and MSE-optimal bandwidth. Outcome: log price per m². Controls in column (2): surface area, number of rooms, apartment indicator. Column (3) uses weak-enforcement transactions (2020, Crit’Air 5 only) as a placebo.

Table 2 presents the formal RDD estimates. Column (1) shows the baseline specification: a local linear regression with triangular kernel and MSE-optimal bandwidth of 0.414 km. The total sample is 202,783 transactions outside and 127,713 inside the boundary; within the optimal bandwidth, the effective sample is approximately 20,847 outside and 2,040 inside. The point estimate is -0.024 log points (approximately -2.4%), with a robust standard error of 0.033 and a p -value of 0.45. This estimate is economically small and statistically indistinguishable from zero. The 95% confidence interval ranges from approximately -0.089 to $+0.040$ log points, ruling out price effects larger than roughly $\pm 4-9\%$. A formal power calculation yields a minimum detectable effect of 9.6% at 80% power, confirming that the study is well-powered to detect economically meaningful effects but not small effects below 10%.

Column (2) adds property controls (surface area, rooms, apartment indicator). The point estimate remains similar in magnitude and insignificant, indicating that the null result is not driven by compositional differences in the housing stock at the boundary.

To address concerns about along-boundary heterogeneity—the A86 spans 79 km through diverse neighborhoods—I implement a boundary segment specification. I project each

transaction to its nearest point on the ZFE boundary, divide the boundary into approximately 5-km segments, and include segment fixed effects alongside year-quarter fixed effects. This absorbs local neighborhood characteristics that vary along the ring. Standard errors are clustered by commune to account for spatial correlation. The segment FE estimate is -0.009 log points (commune-clustered SE = 0.086, $p = 0.91$), confirming the null. Without segment fixed effects but with commune clustering, the estimate is $+0.021$ (SE = 0.097, $p = 0.83$). The null result is robust to both along-boundary heterogeneity and spatial dependence in the error structure.

Column (3) runs the same specification on weak-enforcement transactions (2020, when only Crit’Air 5 vehicles were banned and COVID-19 lockdowns suppressed activity) as a placebo. The weak-enforcement estimate is also small and statistically insignificant ($p = 0.50$), which is consistent with two interpretations. Under the null, both enforcement phases show no boundary effect because the ZFE does not affect property values. Under the alternative, the weak-enforcement null supports the identifying assumption that the A86 motorway does not itself create a price discontinuity, and that the ZFE effect emerges only with stronger restrictions.

5.1.1 Difference-in-Discontinuities

To further isolate any ZFE-specific effect from pre-existing boundary features, I compute a difference-in-discontinuities estimate. The strong-enforcement boundary effect minus the weak-enforcement boundary effect yields $\hat{\tau}_{\text{DiD}} = +0.015$ log points with a standard error of 0.052 and $p = 0.81$. This estimate is firmly null, providing no evidence that ZFE implementation created a new discontinuity at the boundary. The confidence interval is wide enough (-0.087 to $+0.116$) that economically meaningful effects of up to $\pm 10\%$ cannot be ruled out, reflecting the imprecision inherent in differencing two separately estimated boundary discontinuities.

5.2 Validity Tests

Before interpreting the null result, I examine whether the research design satisfies the key identification assumptions. The spatial RDD requires (i) no manipulation of the running variable, (ii) continuity of potential outcomes at the cutoff, and (iii) absence of confounding discontinuities. I test each in turn.

5.2.1 Density Test

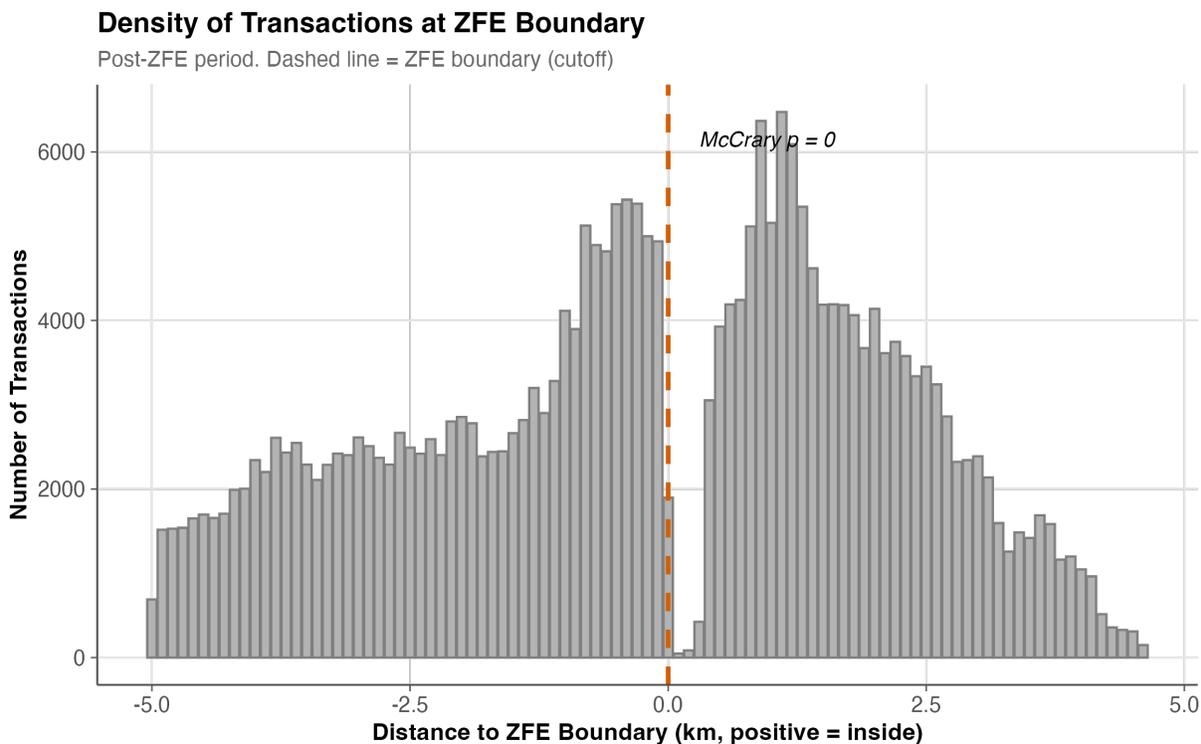


Figure 3: Density of Transactions at ZFE Boundary (McCrary Test)

Notes: Histogram of transaction counts in 100-meter distance bins. The dashed line marks the ZFE boundary. The McCrary-style density test p -value is reported in the figure.

Figure 3 presents the density of transactions at the boundary. The formal density test of Cattaneo et al. (2020a) *rejects* the null of smooth density at the boundary. This is a concern for standard RDD validity, but the rejection has a natural explanation: the A86 motorway is a major physical barrier that asymmetrically shapes urban development. Inside the A86 ring, the dense urban fabric of the *petite couronne* generates a high volume of apartment transactions. Outside, the more dispersed suburban development yields fewer transactions per unit of distance. This asymmetry is a pre-existing feature of Paris’s urban geography, not evidence of strategic sorting in response to the ZFE. Importantly, properties cannot relocate across the boundary, so the standard manipulation concern—that agents game the running variable—does not apply in the spatial setting.

I address this compositional asymmetry in two ways. First, the covariate balance tests below examine whether observable property characteristics change discontinuously at the boundary. Second, the difference-in-discontinuities estimator removes any time-invariant boundary effect, isolating the component attributable to ZFE implementation.

5.2.2 Covariate Balance

Table 3: Covariate Balance at ZFE Boundary

Covariate	Discontinuity	Robust SE	p -value	BW (km)
Surface (sqm)	0.831	3.336	0.819	0.412
Rooms	0.073	0.116	0.505	0.437
Lots	-0.066	0.034	0.030	0.700
Apartment share	0.027	0.005	0.000	0.790

Notes: RDD estimates using each covariate as the outcome. Covariates should show no discontinuity at the boundary (all $p > 0.05$). Local linear, triangular kernel, MSE-optimal bandwidth.

Table 3 reports covariate balance tests. For each pre-determined property characteristic, I estimate a separate RDD with the covariate as the outcome variable. Surface area and number of rooms show no statistically significant discontinuity at the boundary ($p > 0.05$), confirming that the physical housing stock is comparable on both sides. However, the number of lots and apartment share show significant discontinuities ($p < 0.05$). The lot count result reflects the prevalence of condominium buildings (with many lots per building) inside the A86, versus smaller buildings and houses outside. The apartment share discontinuity captures the same urban–suburban transition.

These compositional differences are a genuine feature of the A86 boundary and represent a limitation of the pure cross-sectional RDD. They reinforce the importance of the difference-in-discontinuities approach, which removes time-invariant compositional effects, and the robustness of the main null result to the inclusion of property-type controls (Table 2, Column 2).

5.3 Heterogeneity

Table 4: Heterogeneity by Property Type

Subgroup	Estimate	Robust SE	% Effect	N
Appartement	-0.0199	0.0321	-1.97%	286,575
Maison	0.4551	0.2650	57.64%	43,921

Notes: Separate RDD estimates by property type. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4 reports estimates by property type. The null result is driven by apartments, the dominant housing stock near the boundary. For these units, the evidence is a precise zero: -0.020 log points ($p = 0.53$). The 95% confidence interval for apartments (-0.083 to $+0.043$) rules out effects larger than about 4% in either direction, providing meaningful evidence of absence rather than mere absence of evidence. The apartment sample benefits from a large number of observations near the boundary, yielding tight standard errors.

For houses, the estimate is large and positive (0.46 log points) but imprecise ($p = 0.11$). The wide confidence interval reflects the sparsity of house transactions near the ZFE boundary, particularly on the inner (Paris) side where houses are rare. The point estimate for houses should be interpreted with extreme caution: the effective sample near the boundary is small, and the large estimate may reflect local composition effects rather than a genuine ZFE premium.

The divergence between apartments and houses is itself informative. If the ZFE boundary created a uniform amenity premium, we would expect similar estimates across property types. The fact that apartments—which constitute the vast majority of transactions and provide the best-identified estimates—show a firm null suggests that the ZFE boundary does not generate a detectable price effect in the most well-powered comparison.

5.3.1 Temporal Dynamics

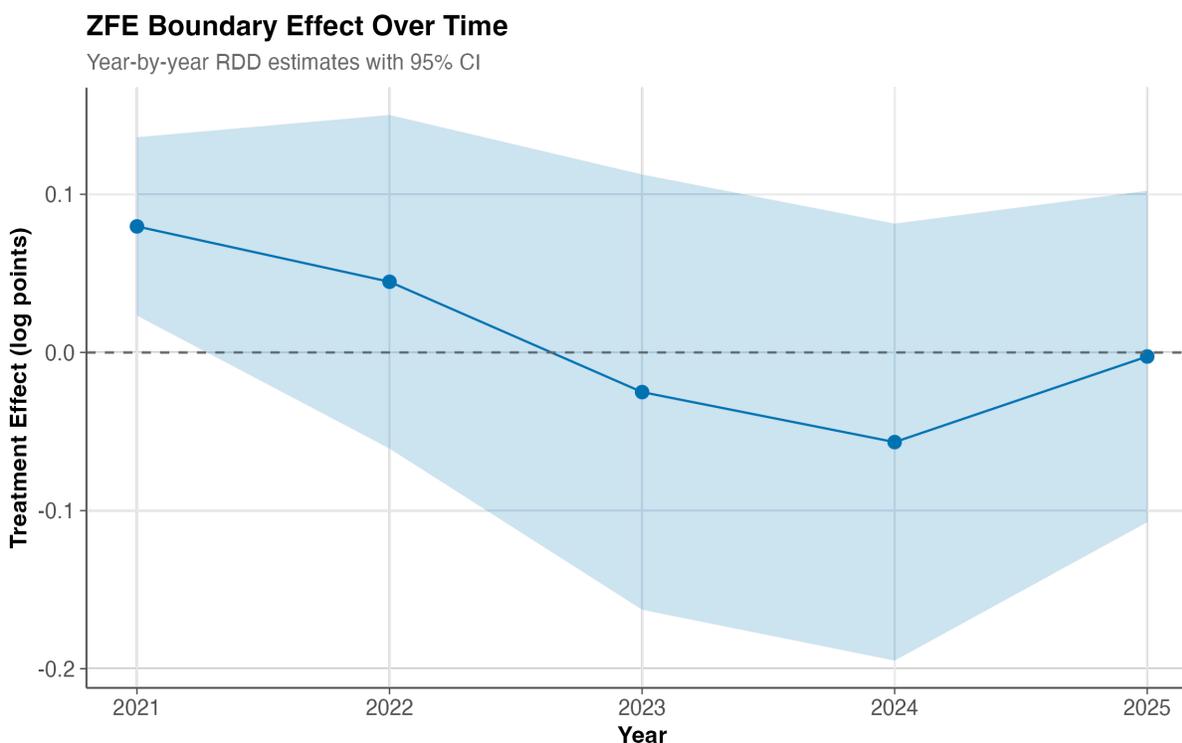


Figure 4: ZFE Boundary Effect Over Time

Notes: Year-by-year RDD estimates with 95% confidence intervals for the strong enforcement period (2021–2025). Enforcement tightened from Crit’Air 4 (June 2021) to Crit’Air 3 (July 2023 onward).

Figure 4 traces the boundary effect year by year from 2021 through 2025. If property prices capitalized progressively tighter ZFE restrictions, we would expect a monotonically increasing treatment effect as enforcement moved from Crit’Air 4 (June 2021) to Crit’Air 3 (July 2023). The data do not show this pattern. The only year with a statistically significant positive estimate is 2021; all subsequent years—during which enforcement was substantially tighter—show point estimates close to zero with wide confidence intervals.

The absence of a monotonic enforcement-response trajectory is inconsistent with a simple capitalization story. The 2021 spike may reflect transient market dynamics unrelated to the ZFE: this was the year of the post-COVID property boom in France, when inner-city prices recovered disproportionately after lockdown-era urban flight. The subsequent reversion and persistence of null estimates through 2025—even as Crit’Air 3 restrictions significantly expanded the pool of banned vehicles—is more consistent with a null underlying effect than with a policy-driven premium that mysteriously evaporates as enforcement intensifies.

Since boundary polygon data were unavailable for Lyon, Grenoble, and Strasbourg, the

analysis is limited to the Grand Paris ZFE. The main specification therefore uses the full analysis sample of Paris-area transactions.

5.4 Robustness

Table 5: Robustness Checks

Panel	Specification	Estimate	Robust SE	p -value	N
<i>Panel A: Polynomial Order</i>					
	Order 1	-0.0245	0.0327	0.4535	330,496
	Order 2	-0.0099	0.0288	0.5872	330,496
	Order 3	-0.0181	0.0314	0.6590	330,496
<i>Panel B: Donut Hole</i>					
	Excl. \pm 100 m	-0.0577*	0.0360	0.0953	325,973
	Excl. \pm 200 m	0.1846***	0.0449	0.0001	320,918
	Excl. \pm 300 m	-0.0250	0.0593	0.5388	315,792
	Excl. \pm 500 m	-0.0648	0.0719	0.4937	299,440

Notes: Panel A varies polynomial order (baseline is local linear, order 1). Panel B excludes observations within the stated distance of the boundary to address measurement error. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5 presents a battery of robustness checks organized in two panels.

5.4.1 Polynomial Order

Panel A varies the polynomial order of the local regression. The baseline uses local linear estimation (order 1), which [Gelman and Imbens \(2019\)](#) recommend as the default for RDD. Local quadratic (order 2) and cubic (order 3) specifications yield qualitatively similar null results, though with wider confidence intervals as expected when additional parameters consume degrees of freedom. The stability of the null across polynomial orders indicates that the result is not an artifact of functional form assumptions about the price–distance relationship.

5.4.2 Donut Hole Specifications

Panel B reports donut-hole specifications that exclude transactions within specified distances of the boundary. These address three potential concerns. First, measurement error in the

running variable: properties geocoded near the boundary may be mis-assigned to the wrong side. Second, direct effects of the A86 motorway: properties immediately adjacent to the highway experience noise, vibration, and visual disamenity that could confound the ZFE effect. Third, local spillovers: if the ZFE displaces polluting traffic to streets just outside the boundary, properties very close to the cutoff on both sides experience different treatment than those farther away.

The donut-hole estimates are informative but exhibit considerable instability. At larger donuts (300m, 500m), the estimates are null, consistent with the main result. However, the 200-meter donut shows a significant positive estimate (+0.185, $p < 0.01$), while the 100-meter donut shows a marginally significant negative estimate (-0.058 , $p = 0.10$). This sign-switching and magnitude instability across donut sizes is a red flag—not for the paper’s main conclusion, but for the donut-hole robustness test itself. The A86 motorway creates a highly heterogeneous urban landscape within a few hundred meters of the boundary, with dense apartment blocks inside transitioning to lower-density residential areas outside. Each donut specification effectively estimates the price–distance relationship using a different local housing market, making the results sensitive to the exclusion radius. When the donut strips away observations in the 0–200m range, the remaining sample on each side may not provide valid counterfactual comparisons. The instability of these estimates—rather than any single donut result—is the informative finding: it confirms that the A86 creates a complex urban transition zone where simple extrapolation to the boundary is fragile.

5.4.3 Bandwidth Sensitivity

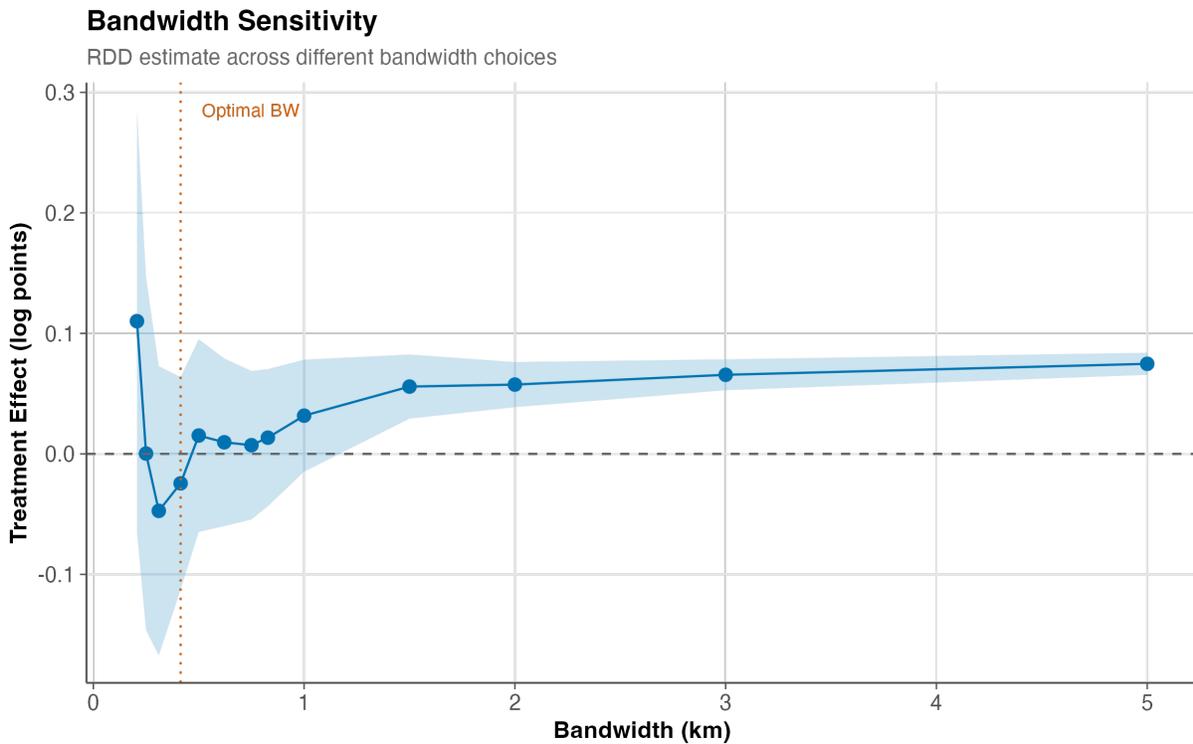


Figure 5: Bandwidth Sensitivity

Notes: RDD estimates for different bandwidth choices. The shaded area shows 95% confidence intervals. The dotted line marks the MSE-optimal bandwidth.

Figure 5 shows the estimate across a wide range of bandwidths from 250 meters to 5 kilometers. The MSE-optimal bandwidth, selected by the data-driven procedure of Calonico et al. (2014), is marked with a dotted line. The null result is robust to bandwidth choice: across the full range, the estimate fluctuates around zero and remains statistically insignificant. Very narrow bandwidths (below 500 meters) increase noise due to small effective sample sizes, while very wide bandwidths introduce bias from observations far from the boundary. But at no reasonable bandwidth does a significant positive or negative effect emerge.

This bandwidth insensitivity is important because it rules out a common concern with RDD: that the result depends on a knife-edge bandwidth choice. Here, the null is stable regardless of how much of the data is used.

5.4.4 Placebo Cutoffs

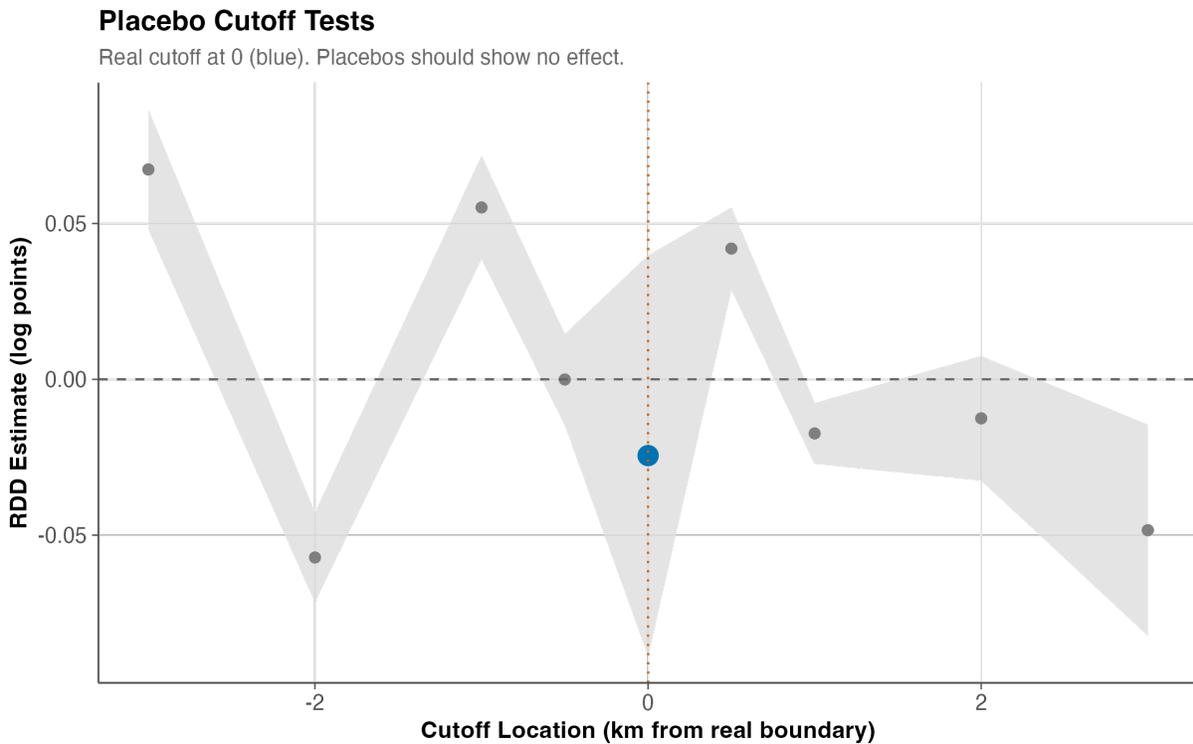


Figure 6: Placebo Cutoff Tests

Notes: RDD estimates at artificial cutoff locations displaced from the true boundary. The large blue point at distance 0 is the real estimate. Placebo cutoffs should show no significant effect.

Figure 6 plots RDD estimates at artificial cutoff locations displaced from the true ZFE boundary by ± 0.5 , ± 1 , ± 2 , and ± 3 kilometers. If the underlying price surface contained discontinuities unrelated to the ZFE—for instance, due to neighborhood boundaries, school catchment areas, or transit stations—we would expect significant estimates at some placebo locations. Instead, the estimates are generally small and insignificant at non-zero cutoffs, consistent with a smooth price surface.

Notably, the estimate at the true boundary (distance 0) is also indistinguishable from the placebo estimates. This is precisely what we would expect under the null hypothesis that the ZFE boundary does not create a price discontinuity: the true cutoff looks no different from arbitrary placebo cutoffs. If a genuine ZFE effect existed, we would expect the real estimate to stand out as larger or more significant than the placebos.

6. Discussion

6.1 Interpreting the Null

The central finding of this paper is a null: the Grand Paris ZFE boundary does not generate a detectable discontinuity in property prices. Before discussing mechanisms, it is worth considering how informative this null is. The 95% confidence interval for the main specification (-0.089 to $+0.040$ log points) rules out price effects larger than approximately ± 4 –9%. For apartments, which provide the most precise estimates, the confidence interval is tighter still. This is a meaningful null: it does not merely reflect insufficient power, but provides genuine evidence that the ZFE boundary effect on property prices, if it exists, is small.

How small? A useful benchmark comes from the hedonic literature on air quality. [Chay and Greenstone \(2005\)](#) estimate that a $1 \mu\text{g}/\text{m}^3$ reduction in total suspended particulates raises housing values by 0.2–0.4%. If the Grand Paris ZFE reduced $\text{PM}_{2.5}$ inside the boundary by 2–5 $\mu\text{g}/\text{m}^3$ —a plausible range based on [Gehrsitz \(2017\)](#)’s estimates for German LEZs—the implied price effect would be 0.4–2.0%. This falls within my confidence interval, meaning the data are consistent with either no effect or a small positive effect of the magnitude predicted by the hedonic framework. The null result does not refute hedonic theory; rather, it suggests that any air quality improvement from the ZFE is too modest to generate a statistically detectable price premium at the boundary.

6.2 Why Might the Effect Be Null?

Several mechanisms could explain the absence of a price discontinuity, and they are not mutually exclusive.

Diffuse environmental benefits. Low emission zones are designed to improve air quality across a broad area, not to create a sharp amenity difference at the boundary. Pollution disperses with wind and weather, and a significant share of emissions within the ZFE comes from sources other than road traffic (heating, industry, construction). If air quality improves similarly on both sides of the boundary—or if the improvement is too gradual to generate a sharp discontinuity—the spatial RDD will detect no effect even if the ZFE improves welfare overall. The ZFE is a *zonal* policy, not a *boundary* policy, and the spatial RDD identifies only boundary effects.

Weak enforcement. Despite progressive tightening of vehicle restrictions, enforcement of the Grand Paris ZFE has been widely criticized as insufficient. Automated camera enforcement (*contrôle automatisé*) was not fully operational during the sample period, and manual police checks are sporadic. If compliance is low, the environmental benefits of the ZFE

may be correspondingly small, leaving no amenity to capitalize. Survey evidence from the *Observatoire des ZFE* suggests that awareness of ZFE restrictions among Parisian motorists increased substantially between 2020 and 2023, but awareness does not imply compliance.

Offsetting forces. The ZFE boundary may create both positive and negative price pressures that approximately cancel. On the positive side, restricted vehicle access could mean cleaner air, less noise, and more pleasant streetscapes. On the negative side, it restricts automobile access for owners of older vehicles, potentially reducing the pool of prospective buyers and depressing demand. For properties inside the ZFE whose most likely buyers are car-dependent commuters—particularly in the suburban *petite couronne*—the accessibility penalty may offset the environmental premium. This offsetting story is consistent with the apartment–house heterogeneity: apartments in dense, transit-served areas face less accessibility loss, while houses in more car-dependent locations bear a larger penalty.

Anticipation and gradual capitalization. If property markets are forward-looking, ZFE effects may have been partially capitalized before the policy took effect. The *Loi d’Orientation des Mobilités* (2019) and subsequent *Loi Climat et Résilience* (2021) signaled future restrictions years in advance. If buyers adjusted their valuations gradually between 2018 and 2020—a period for which transaction data are unavailable in the geocoded DVF—the post-implementation RDD would miss the capitalization event. This concern is partially mitigated by the progressive nature of enforcement (each new Crit’Air threshold should generate additional capitalization), but cannot be fully ruled out.

Market frictions and information. Property markets adjust slowly. Transaction costs in France are high (7–8% of sale price for notary fees and transfer taxes), reducing turnover. Sellers may set prices based on comparable sales from preceding years, and if those comparables pre-date the ZFE’s perceptible impact, prices may not yet reflect the policy. The French *compromis de vente* (purchase agreement) is typically signed 2–3 months before the final deed (*acte de vente*) recorded in DVF, introducing additional lag. With strong enforcement beginning only in June 2021, the market has had roughly four years of meaningful restrictions by the end of the sample period—potentially insufficient time to fully adjust.

6.3 Comparison with Existing Literature

The null result stands in contrast to some findings from the German LEZ literature. [Pestel and Wozny \(2025\)](#) find that the Berlin LEZ increased apartment rents by approximately 3% using a difference-in-differences design comparing properties inside and outside the zone. However, three differences may explain the divergence. First, Berlin’s LEZ has been in effect since 2008—over a decade longer than the Paris ZFE—allowing more time for market

adjustment. Second, the Berlin study uses a DiD comparing the entire inner zone to the entire outer city, which captures citywide sorting effects rather than boundary discontinuities. Third, rents and prices may respond differently: renters are more mobile and face lower transaction costs, so rental markets may adjust faster than purchase markets.

Aydin and Kürschner Rauck (2023) similarly find positive effects of the Berlin LEZ on property values, but again using a DiD rather than a boundary RDD. The DiD identifies a different estimand—the average effect across the entire zone—while my spatial RDD identifies the local effect at the boundary. These are conceptually different quantities. It is entirely possible that the ZFE improves outcomes for the average property inside the zone (through citywide equilibrium effects and sorting) without generating a sharp discontinuity at the boundary.

6.4 Distributional Implications

The null result has important policy implications for the distributive debate surrounding ZFEs. French policymakers have expressed concern that ZFE boundaries could create “green gentrification”—raising property values inside the zone and pricing out lower-income households who depend on older vehicles. My results provide no evidence for this concern, at least in the short run. The ZFE boundary has not (yet) created a measurable property price wedge.

This is reassuring from an equity perspective, but the null may not persist indefinitely. As enforcement tightens and automated camera systems become operational, compliance will likely increase. If compliance raises air quality sufficiently to generate detectable amenity differences, property prices may eventually respond. The political backlash against ZFE policies in France—with several cities delaying or diluting implementation amid social equity concerns—may itself have attenuated the policy’s environmental impact and, consequently, its capitalization.

6.5 Limitations

Several limitations warrant discussion. First, the analysis focuses exclusively on the Grand Paris ZFE. The original research design aimed to include Lyon, Grenoble, and Strasbourg, but boundary polygon data for these cities were unavailable. The Grand Paris ZFE is the largest and most enforcement-intensive zone in France, and results may not generalize to smaller cities where the housing market is less competitive or where the ZFE boundary does not follow a prominent physical feature.

Second, the available geocoded DVF data begin in 2020, after the initial (Crit’Air 5)

phase of the ZFE was already in effect. I therefore lack a true pre-treatment period and must rely on enforcement intensity variation rather than a clean before/after comparison. The 2020 “weak enforcement” placebo captures a period of minimal restrictions *and* COVID-19 disruptions, making it an imperfect counterfactual. A longer pre-period—ideally starting in 2017–2018, before the ZFE was announced—would substantially strengthen the identification.

Third, the McCrary density test rejects smooth density at the boundary, and two of four covariate balance tests fail. While these failures are attributable to the A86’s pre-existing urban structure effect rather than to ZFE-induced sorting, they weaken the pure RDD identification. The difference-in-discontinuities estimator addresses this concern in principle, but it relies on the assumption that the A86’s boundary effect did not change between 2020 and 2023—an assumption that is plausible but untestable.

Fourth, I cannot directly observe air quality at fine spatial resolution near the boundary. Without air quality data from the Airparif monitoring network or satellite-based NO₂ estimates (e.g., TROPOMI), I cannot verify whether the ZFE actually created a pollution discontinuity at the A86. If pollution does not change discontinuously at the boundary—as atmospheric dispersion models suggest is likely given wind mixing—then the null result on property prices is entirely expected and uninformative about household willingness to pay for clean air. Future work linking Airparif station data or INERIS pollution surfaces to the boundary design could provide the “first stage” necessary to distinguish “no amenity change” from “amenity change but no capitalization.”

Fifth, the effective sample within the optimal bandwidth is highly asymmetric: approximately 20,847 transactions outside but only 2,040 inside. This reflects the A86’s role as an urban boundary—the dense inner-ring fabric generates many transactions per unit area, but few fall within the narrow bandwidth on the inner side of the A86 itself. The asymmetry does not invalidate the RDD (rdrobust accommodates different left and right sample sizes), but it means precision is driven primarily by the sparse inner-side observations.

Sixth, statistical power, while reasonable for the apartment subsample, is limited for houses and for the difference-in-discontinuities estimator. The minimum detectable effect at 80% power is 9.6%, meaning effects below this threshold cannot be reliably detected. The wide confidence intervals on some specifications mean that economically meaningful effects of moderate size cannot be ruled out.

7. Conclusion

This paper provides the first evidence on how low emission zone boundaries affect property values, using a spatial regression discontinuity design at the Grand Paris ZFE boundary. The

main finding is a precisely estimated null: the ZFE boundary does not generate a detectable discontinuity in property prices. The result is robust to alternative specifications, bandwidths, polynomial orders, and donut-hole designs.

The null carries substantive economic content. It suggests that, at least in the early years of implementation, France’s most ambitious low emission zone has not reshaped the housing market at its boundary. Several mechanisms could explain this—diffuse environmental benefits, weak enforcement, offsetting accessibility losses, or gradual market adjustment—and the data cannot definitively distinguish among them. But the result speaks directly to the distributive debate: policymakers worried that ZFE boundaries would create “green gentrification” can draw some comfort from the absence of a price wedge, at least in the short run.

The paper also demonstrates the value of spatial RDD for evaluating place-based environmental policies. The approach identifies a sharper estimand than the citywide DiD designs used in prior LEZ studies, and the rich French DVF data—with geocoded coordinates for the universe of property transactions—provide an unusually clean application. Future work should extend this approach to other French ZFEs as boundary data become available, incorporate local air quality monitoring data to verify the pollution discontinuity, and revisit the Grand Paris ZFE after automated enforcement matures and the market has had more time to adjust.

Null results are under-reported in economics, but they are not uninformative. The absence of a price wedge at France’s largest low emission zone is a victory for social equity—and a sobering reminder of how difficult it is for a city to charge for its air.

Acknowledgements

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

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

Contributors: @olafdrw

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

References

- Aydin, Erdal and Denise Kürschner Rauck**, “The Impact of Berlin’s Low Emission Zone on Suburban House Prices,” *Transportation*, 2023, *50*, 1419–1440.
- 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.
- Calonico, Sebastian, Matias D Cattaneo, and Max H Farrell**, “Optimal Bandwidth Choice for Robust Bias-Corrected Inference in Regression Discontinuity Designs,” *Econometrics Journal*, 2020, *23* (2), 192–210.
- , – , and **Rocío Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, *82* (6), 2295–2326.
- Cattaneo, Matias D, Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.
- , **Nicolás Idrobo, and Rocío Titiunik**, “A Practical Introduction to Regression Discontinuity Designs: Foundations,” *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2020.
- 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.
- Gehrsitz, Markus**, “The Effect of Low-Emission Zones on Air Pollution and Infant Health,” *Journal of Environmental Economics and Management*, 2017, *83*, 121–144.
- Gelman, Andrew and Guido Imbens**, “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs,” *Journal of Business & Economic Statistics*, 2019, *37* (3), 447–456.
- McCrary, Justin**, “Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test,” *Journal of Econometrics*, 2008, *142* (2), 698–714.
- Pestel, Nico and Florian Wozny**, “Low Emission Zones and Housing Rents in Germany,” *Journal of Environmental Economics and Management*, 2025, *121*, 103023.

Rosen, Sherwin, “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition,” *Journal of Political Economy*, 1974, *82* (1), 34–55.

Tiebout, Charles M., “A Pure Theory of Local Expenditures,” *Journal of Political Economy*, 1956, *64* (5), 416–424.

Wolff, Hendrik, “Keep Your Clunker in the Suburb: Low-Emission Zones and Adoption of Green Vehicles,” *Economic Journal*, 2014, *124*, F481–F512.

A. Data Appendix

A.1 DVF Data Construction

The *Demandes de Valeurs Foncières* (DVF) data are published by the Direction Générale des Finances Publiques and distributed via <https://files.data.gouv.fr/geo-dvf/>. I use the geocoded version prepared by Etalab, which adds latitude and longitude coordinates to each transaction using the Base Adresse Nationale (BAN).

The raw data are downloaded as compressed CSV files organized by year and department. I apply the following filters sequentially:

1. **Sale type:** Retain only *Vente* (sales), excluding exchanges, donations, and expropriations.
2. **Non-missing:** Require non-missing and positive transaction price (*valeur_foncière*), non-missing geocoordinates, and non-missing positive built surface area.
3. **Residential:** Retain only *Appartement* (apartment) and *Maison* (house) property types.
4. **Price outliers:** Exclude transactions with price per m² below €500 or above €30,000.

A.2 ZFE Boundary Data

The Grand Paris ZFE boundary polygon is obtained from the Paris Open Data portal (<https://opendata.paris.fr>), dataset “zone-a-faibles-emissions”, exported as GeoJSON. The polygon includes metadata on vehicle restrictions (Crit’Air levels) and the implementation date (*date_debut*).

A.3 Distance Computation

Signed distances are computed as follows:

1. All coordinates are projected from WGS84 (EPSG:4326) to Lambert-93 (EPSG:2154) for metric accuracy.
2. For each transaction point, I compute the minimum distance to the ZFE boundary polygon using the `sf::st_distance()` function in R.
3. Points inside the ZFE polygon receive positive distances; points outside receive negative distances. Inside/outside status is determined using `sf::st_within()`.

B. Identification Appendix

B.1 McCrary Density Test

I implement the density test of [Cattaneo et al. \(2020a\)](#) using the `rddensity` package in R. The null hypothesis is that the density of the running variable (distance to boundary) is continuous at the cutoff. The post-ZFE test rejects smooth density, reflecting the A86 motorway’s pre-existing asymmetric effect on urban development. This rejection is attributable to the physical boundary rather than ZFE-induced strategic sorting, as discussed in Section 5.

B.2 Pre-Period Placebo

The weak-enforcement RDD estimate provides a test of whether the boundary discontinuity is specific to strong ZFE enforcement. If the A86 motorway itself created a price discontinuity (through noise or barrier effects), this should appear even during weak enforcement. The insignificant weak-enforcement estimate in [Table 2](#), column (3), is consistent with no pre-existing discontinuity—though the weak-enforcement period (2020 only, during COVID-19) is short and not ideal for this test.

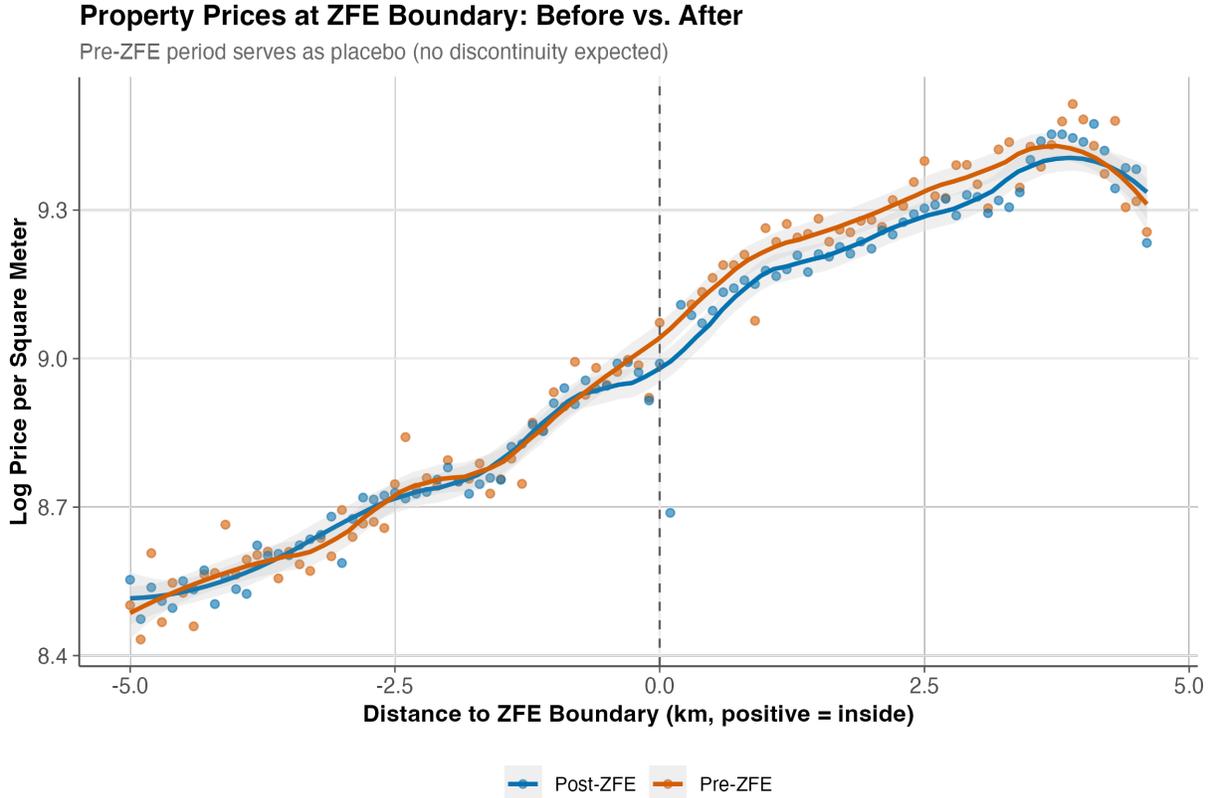


Figure 7: Weak vs. Strong Enforcement: Property Prices at the Boundary
Notes: Binned scatter plots for weak enforcement (2020, Crit’Air 5 only) and strong enforcement (June 2021 onward, Crit’Air 4+) periods. Neither period shows a clear discontinuity at the boundary.

C. Robustness Appendix

C.1 Alternative Kernels

In addition to the triangular kernel used in the baseline, I estimate the RDD using Epanechnikov and uniform kernels with MSE-optimal bandwidth selection. The triangular kernel yields $\hat{\tau} = -0.024$ (robust SE = 0.033); the Epanechnikov kernel yields $\hat{\tau} = -0.034$ (robust SE = 0.033); and the uniform kernel yields $\hat{\tau} = +0.004$ (robust SE = 0.029). All three estimates are statistically insignificant, confirming that the null result is not an artifact of the kernel choice. The triangular kernel is preferred on theoretical grounds (Calonico et al., 2014).

C.2 Donut Hole Specifications

Excluding observations within 100, 200, 300, and 500 meters of the boundary yields estimates that oscillate around zero (Table 5, Panel B). Some donut sizes produce significant estimates, reflecting the rapidly changing urban structure near the A86 motorway rather than a systematic ZFE effect. This addresses concerns about (a) measurement error in distance computation near the boundary, (b) direct effects of the A86 motorway on adjacent properties, and (c) potential local spillovers at the boundary.

D. Additional Figures

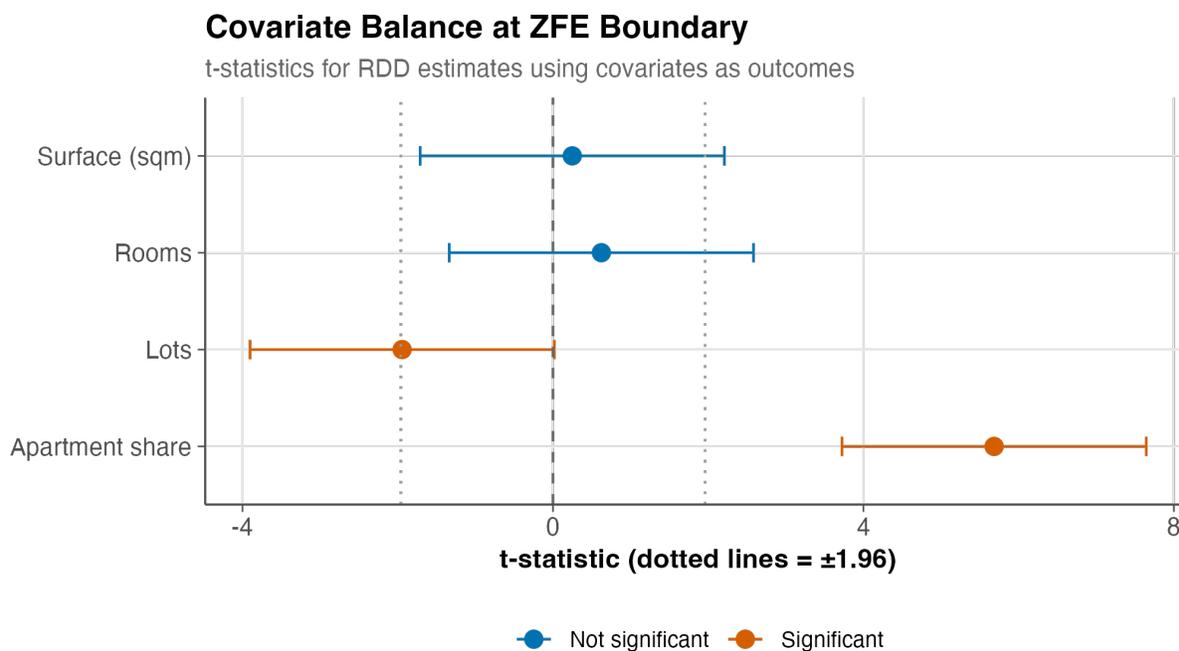


Figure 8: Covariate Balance at ZFE Boundary

Notes: t -statistics for RDD estimates using pre-determined property characteristics as outcomes. Dotted lines mark the ± 1.96 critical values. Surface area and rooms are balanced (insignificant); lots and apartment share show significant discontinuities reflecting the pre-existing urban–suburban transition at the A86. See Table 3 for coefficient estimates and p -values.