The Stigma of Priority: Education Priority Zone Labels and Housing Prices in France

APEP Autonomous Research* @olafdrw

March 4, 2026

Abstract

Do education priority zone labels affect property values? I study price gradients between France's REP and non-REP collèges using 1.7 million geocoded property transactions (2020–2024). The naïve correlation between REP proximity and prices is −14.2 percent, but progressively adding distance controls, property characteristics, département and commune fixed effects reduces the estimate to exactly zero (−0.000, SE = 0.008). The gradient declined from 7.7 percent in 2020 to 2.4 percent in 2024 and reverses sign outside the Paris region. In areas with abundant private schools, the gradient reverses (−2.1 percent), consistent with private schools neutralizing sorting incentives. REP labels do not impose a measurable stigma tax; the observed price gap is entirely attributable to geographic sorting.

JEL Codes: H75, I24, R21, R31

Keywords: education priority zones, housing prices, school quality, boundary discontinuity, France

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

1. Introduction

Do government labels intended to help disadvantaged schools instead stigmatize the neighborhoods they serve? France's *Réseau d'Éducation Prioritaire* (REP) program channels additional resources—smaller classes, supplementary teachers, bonus pay—to schools in disadvantaged areas. While such interventions aim to reduce educational inequality, the very act of labeling a school as "priority" may signal neighborhood deprivation to prospective homebuyers, depressing property values. Understanding whether and how these labels capitalize into housing prices is central to evaluating place-based education policies, because any stigma effect creates a negative feedback loop: lower property values reduce the local tax base, further disadvantaging the labeled neighborhoods the policy was designed to help.

A large literature, beginning with Black (1999), documents that school quality capitalizes into property values. Families pay substantial premiums to access better-rated schools, with boundary discontinuity designs revealing sharp price gradients at attendance zone borders in the United States (Bayer et al., 2007; Kane et al., 2006; Cellini et al., 2010), the United Kingdom (Gibbons et al., 2013; Cheshire and Sheppard, 2004; Machin, 2011), and France (Fack and Grenet, 2010). The typical finding is that a one-standard-deviation improvement in test scores is associated with a 2–5 percent increase in nearby property prices. Yet most of this literature examines school *performance* signals—test scores, school grades, rankings—rather than government *policy labels* that designate schools as serving disadvantaged populations. The distinction is important: a performance signal provides information about outcomes, while a priority label conveys information about the *composition* of the student body and the *character* of the neighborhood. A priority label may depress property values even when the policy it triggers actually improves school quality, because the label's negative signal about neighborhood composition may dominate any positive signal about resource provision.

This paper examines the housing market consequences of France's REP designation by studying price gradients at the equidistance loci between REP and non-REP *collèges*. I construct a novel dataset linking 4.6 million geocoded property transactions from France's comprehensive *Demandes de Valeurs Foncières* (DVF) database with the universe of public *collèges* (lower secondary schools) and their REP status. For each transaction, I compute the distance to the nearest REP and nearest non-REP collège, constructing a signed distance running variable that equals zero at the locus of points equidistant from the two school types. Because actual *carte scolaire* catchment boundaries are not publicly available, this equidistance locus serves as a proxy that approximates—but does not perfectly correspond to—the administrative assignment boundary. Comparing properties on either side of this proxy boundary measures the price gradient associated with relative REP proximity, a composite of any label effect and correlated neighborhood differences.

A nonparametric local polynomial comparison at the equidistance boundary, using the Calonico et al. (2014) bias-corrected estimator with MSE-optimal bandwidth (57 meters), reveals a 5.3 percent price gradient ($z = 12.4$, $p < 0.001$), with properties on the REP side of the equidistance locus selling for *more* than those on the non-REP side. However, standard RDD assumptions are violated in this setting—the McCrary density test rejects and covariates are discontinuous—so this estimate captures the combined effect of the label and correlated neighborhood differences, not a locally randomized treatment effect. This estimate is stable across bandwidth choices from half to double the optimal bandwidth, ranging from 4.8 to 6.6 percent. However, the boundary RDD in this setting measures a composite of the label effect and neighborhood differences correlated with REP designation. Parametric specifications that progressively add distance controls, property characteristics, and département fixed effects tell a revealing story: the naïve correlation is −14.2 percent, adding the distance polynomial flips the sign to +6.3 percent, property controls reduce it to +4.3 percent, département fixed effects bring it to +1.3 percent (insignificant), and commune fixed effects yield exactly zero (−0.000, SE = 0.008). The R^2 rises from 0.01 to 0.61, indicating that geographic sorting at the regional and municipal levels explains nearly all the variation in property prices near REP boundaries.

Several additional findings emerge from the analysis. First, the boundary gap declined substantially over the study period, from 7.7 percent in 2020 to 2.4 percent in 2024 (and 2.0 percent in partial 2025 data covering January–February only), suggesting convergence possibly driven by post-COVID housing market dynamics, the accumulation of positive signals from the *dédoublement* class size reduction policy launched in REP schools since 2017, or evolving perceptions of education priority labels. Second, excluding the Île-de-France region reverses the sign of the boundary gap (−2.7 percent), demonstrating that the baseline finding is driven entirely by the Paris metropolitan area, where REP and non-REP neighborhoods differ most sharply in socioeconomic composition.

Third, and most strikingly, the gap reverses in areas with high private school density. Where the nearest private collège alternative is readily available, the boundary gap flips to −2.1 percent ($p < 0.001$), compared to +2.7 percent in low-density areas. This asymmetry is consistent with private schools functioning as an "escape valve" that neutralizes the residential sorting incentive created by the REP label (Hoxby, 2000; Nechyba, 2006). When families can avoid the REP collège by enrolling in a nearby private school, the incentive to move across the catchment boundary disappears, and the boundary price gradient vanishes.

Fourth, the McCrary density test rejects the null of no sorting at the boundary ($T = 31.3$, $p < 0.001$), and covariates—surface area, number of rooms, apartment share—are discontinuous at the threshold. These violations of standard RDD assumptions are expected in the housing market context: unlike administrative RDD designs where the running variable is difficult to manipulate, families endogenously choose where to live, and the resulting equilibrium features systematic differences across REP boundaries. The identifying variation in this setting is therefore better interpreted as measuring the *gradient* of property values associated with REP proximity, conditional on distance controls, rather than a locally randomized treatment effect. I discuss these identification challenges transparently and argue that the parametric specifications with rich geographic fixed effects provide the most credible estimates.

This paper contributes to several literatures. First, it extends the school quality capitalization literature (Black, 1999; Bayer et al., 2007; Gibbons et al., 2013; Machin, 2011) to the specific question of government *stigma labels*—signals that a school serves disadvantaged populations—as opposed to performance metrics. Fack and Grenet (2010) study Paris school quality capitalization using boundary discontinuities but focus on brevet pass rates rather than REP status, and their analysis is restricted to the city of Paris. I study the REP label nationally and find that its housing market impact is substantially weaker than that of performance signals documented in the prior literature. Figlio and Lucas (2004) and Friesen et al. (2012) examine the housing market effects of school report cards and rankings; I study a different type of signal—an explicit government designation of disadvantage that operates through neighborhood composition rather than school quality information.

Second, it contributes to the literature on France's education priority policy. Bénabou et al. (2009) find that the original ZEP designation had no detectable effect on student outcomes ("much ado about nothing"); my findings suggest a parallel result in housing markets once geographic sorting is controlled. The REP label appears to be neither strongly beneficial (through signaling additional resources) nor strongly harmful (through stigma) to local property values—a genuine null result that is informative about the equilibrium response to place-based education labeling.

Third, I provide new evidence on the role of private school competition in moderating the capitalization of public school characteristics (Hoxby, 2000; Nechyba, 2006; Fack and Grenet, 2010). Fack and Grenet (2010) show that private school availability attenuates the capitalization of public school test scores in Paris; I extend this finding nationally and show that private alternatives specifically neutralize the *stigma* component of the REP boundary gap. The mechanism is that private schools expand the effective choice set facing families, relaxing the tight link between residential location and school assignment that is the foundation of Tiebout sorting (Tiebout, 1956).

Fourth, the paper demonstrates the potential of France's open-access DVF transaction data for spatial economic research, constructing a national boundary discontinuity design at an unprecedented scale of 1.7 million transactions across 6,429 school locations.

The remainder of the paper proceeds as follows. Section 2 describes France's education priority system and the institutional setting. Section 3 presents a simple conceptual framework linking school labels to property prices. Section 4 presents the data sources and sample construction. Section 5 details the empirical strategy and discusses identification. Section 6 presents the main results, robustness checks, and mechanism analysis. Section 7 discusses interpretation and limitations. Section 8 concludes.

2. Institutional Background

2.1 France's Education Priority System

France's education priority policy (*Éducation Prioritaire*) was established in 1981 under Education Minister Alain Savary as one of the first systematic place-based education interventions in Europe. The program's founding principle was "giving more to those who have less" (*donner plus à ceux qui ont moins*), channeling additional teaching resources and funding to schools serving socioeconomically disadvantaged populations. Over four decades, the system has undergone multiple reforms—in 1990, 1998, 2006, and most recently 2015—that have reshaped both the set of designated schools and the intensity of resources they receive.

The original 1981 program created *Zones d'Éducation Prioritaire* (ZEP), geographic areas containing clusters of schools with high concentrations of disadvantaged students. The ZEP designation provided additional teaching positions, supplementary hours, and modest funding for pedagogical projects. Subsequent reforms added additional labels—*Réseau Ambition Réussite* (RAR) in 2006, *Programme Éclair* in 2011—creating a complex layering of designations that confused both schools and families.

The most recent comprehensive reform occurred in 2014–2015, when Education Minister Najat Vallaud-Belkacem reorganized the entire system into two clear tiers: REP (*Réseau d'Éducation Prioritaire*) and REP+ (*Réseau d'Éducation Prioritaire renforcé*). The REP+ tier targets the most disadvantaged schools—roughly the bottom 5 percent of the socioeconomic distribution—and provides the most intensive support. As of 2022–2023, the education priority network comprises approximately 1,100 collèges at the lower secondary level (roughly 15 percent of all public collèges) and their associated primary schools, serving about 1.7 million students nationwide.

2.2 Designation Criteria

The designation criteria are based on a composite social index (*indice de position sociale*, IPS) that aggregates four dimensions of the enrolled student body: parental socioprofessional categories, educational attainment, household income characteristics, and family structure indicators. The IPS is computed at the school level from administrative enrollment records. Schools below certain thresholds on this composite index receive REP or REP+ status, with the REP+ threshold being more restrictive. The designation is reviewed periodically but has remained relatively stable since the 2015 reform.

Importantly, the designation is attached to the school (specifically to the collège that anchors the local network) rather than to a fixed geographic boundary. This means that REP status reflects the *current* student body composition, which in turn reflects the residential catchment area. The circularity—neighborhood composition determines REP status, which may in turn affect neighborhood composition through housing market responses—is a central identification challenge that I address in Section 5.

2.3 The Carte Scolaire and School Assignment

French public school assignment is governed by the *carte scolaire*, a catchment area system established in 1963 that maps each residential address to a designated public collège. Families residing within a catchment area are assigned to the corresponding school by default. While derogation requests (*demandes de dérogation*) allow some families to attend a different public school, acceptance rates are limited—typically 10–20 percent of requests are granted—and vary by *académie* and the availability of places. Medical reasons, sibling enrollment, and proximity to the alternative school receive priority in derogation decisions.

The carte scolaire was briefly relaxed under the Sarkozy presidency (2007–2012), when families in education priority zones were given priority for derogation requests. This "assouplissement" policy was subsequently reversed, and the current system maintains relatively strict catchment boundaries. The practical consequence is that residential location strongly determines public school assignment, creating the preconditions for Tiebout sorting (Tiebout, 1956): families who value school quality have a strong incentive to locate in catchment areas of desirable schools, bidding up property prices near those schools.

The carte scolaire generates sharp boundaries between adjacent catchment areas. Properties separated by a single street may be assigned to different collèges—one REP-designated, the other not. This institutional feature motivates the boundary discontinuity design: comparing properties near these boundaries exploits the fact that nearby households live in similar neighborhoods but face different school assignments. Fack and Grenet (2010) pioneered this approach in Paris, exploiting catchment boundaries within the city to estimate the capitalization of brevet exam pass rates. I extend this approach nationally and focus on the REP label rather than performance metrics.

2.4 Resources in REP/REP+ Schools

REP and REP+ schools receive several categories of additional resources. The flagship policy since 2017 is *dédoublement*—the systematic halving of class sizes in the first two years of primary school (*Cours Préparatoire* and *Cours Élémentaire 1*) in REP+ schools, later extended to REP schools and to the third year (*CE2*). This represents a substantial investment: class sizes in targeted grades fell from roughly 24 students to 12 students, requiring the recruitment of approximately 10,000 additional teachers.

Teachers in REP+ schools receive an annual salary bonus (*indemnité de sujétion*) of approximately 2,312 euros (raised to 5,114 euros in 2023 as part of the "Pacte enseignant" reform), and benefit from 18 half-days of dedicated professional development time per year—time released from classroom instruction for collaborative pedagogical work. REP (non-plus) teachers receive a smaller bonus (1,734 euros, raised to 3,868 euros). REP+ collèges also receive additional teaching hours (*heures supplémentaires*) and support staff including educational assistants, school nurses, and social workers.

Early evaluations by the Direction de l'évaluation, de la prospective et de la performance (DEPP) suggest modest positive effects of the class size reduction on student achievement in the early primary grades, particularly in mathematics (DEPP, 2019). Piketty and Valdenaire (2006) had earlier documented that class size reductions in French ZEP schools improved student performance, consistent with the broader international evidence (Angrist and Lavy, 1999; Krueger, 1999; Fredriksson et al., 2013). However, the REP label simultaneously signals neighborhood disadvantage to prospective homebuyers. The net effect on property values depends on whether families value the additional resources more than they are deterred by the stigma signal—a question this paper addresses empirically.

3. Conceptual Framework

3.1 School Labels in the Hedonic Price Model

I develop a simple framework to organize the empirical analysis. Consider a hedonic housing market (Rosen, 1974) in which the price of property i reflects its structural characteristics

($\mathbf{Z}_i$), neighborhood amenities ($\mathbf{A}_i$), and the characteristics of the assigned school ($\mathbf{S}_i$):

$$\ln p_i = \mathbf{Z}_i'\boldsymbol{\beta} + \mathbf{A}_i'\boldsymbol{\gamma} + \mathbf{S}_i'\boldsymbol{\delta} + \varepsilon_i \tag{1}$$

The school vector $\mathbf{S}_i$ includes both objective quality (q_i, e.g., test scores, teacher qualifications) and the government label ($L_i \in \{0,1\}$, where $L_i = 1$ indicates REP status). The label effect can be decomposed into two channels:

$$\frac{\partial \ln p_i}{\partial L_i} = \underbrace{\frac{\partial \ln p_i}{\partial q_i} \cdot \frac{\partial q_i}{\partial L_i}}_{\text{Resource channel } (+)} + \underbrace{\left.\frac{\partial \ln p_i}{\partial L_i}\right|_{q_i}}_{\text{Stigma channel } (-)} \tag{2}$$

The resource channel is positive: REP designation brings additional teachers, smaller classes, and supplementary funding that improve school quality, which families value. The stigma channel is negative: the REP label signals that the school serves a disadvantaged population, conveying negative information about neighborhood composition that reduces demand. The net effect is ambiguous *a priori*.

3.2 The Role of Private School Alternatives

The hedonic framework yields a clear prediction about the role of private schools. Following Nechyba (2006) and Fack and Grenet (2010), consider two types of families: those constrained to the public system (who must live in the catchment area of their desired public school) and those with access to private alternatives (who can choose between public and private schools regardless of location).

For constrained families, the REP label directly affects the demand for housing in the catchment area: families who wish to avoid the REP school must physically relocate, creating a sorting premium at the catchment boundary. For unconstrained families, the REP label is less consequential: they can enroll their children in a nearby private school without moving, so the label does not generate residential sorting pressure.

This yields a testable prediction:

> *Prediction:* The REP boundary gap should be larger in areas with low private school density (where most families are constrained to the public system) and smaller in areas with high private school density (where families have exit options).

If private schools fully neutralize the sorting incentive, the boundary gap should be zero in high-density areas, regardless of whether the stigma channel is large or small. The key insight is that private schools do not need to eliminate the *perception* of stigma—they need only eliminate the *residential sorting response* to that perception.

3.3 Equilibrium Implications

In the Bénabou (1993) model of cities with local public goods, school quality and neighborhood composition are jointly determined in equilibrium. REP designation creates a discontinuity in the "label" component of school characteristics, but the resulting sorting equilibrium may feature smooth transitions in neighborhood quality across the boundary. This is important for interpreting the RDD estimates: even if the label effect is sharp, the equilibrium response—families sorting in response to the label—smooths the observable price function.

The implication is that the boundary RDD measures the *equilibrium* price gradient at the REP threshold, which combines the direct label effect with the indirect effects of sorting. The parametric specifications with geographic controls attempt to absorb the sorting component, yielding an estimate closer to the direct label effect. However, perfect decomposition is not possible without exogenous variation in the label conditional on neighborhood characteristics.

4. Data

4.1 Property Transactions: DVF

The primary outcome data come from the *Demandes de Valeurs Foncières* (DVF) database, France's comprehensive register of property transactions. Established by France's tax authority (Direction générale des Finances publiques, DGFiP), DVF records every property transaction in mainland France with the associated sale price. Since April 2019, DVF data are freely available through `data.gouv.fr` in geocoded form, making France one of the few countries with a comprehensive, open-access register of property prices with precise geographic coordinates.

Each DVF record includes: the transaction price, mutation date, property type (apartment, house, commercial, or land), total built surface area, number of principal rooms, commune and département codes, and geographic coordinates precise to approximately 10 meters. The dataset does not include property condition, renovation status, energy performance rating (DPE), or detailed neighborhood characteristics such as crime rates or local amenities.

I use the consolidated geocoded DVF dataset covering transactions from January 2020 through February 2025, comprising 20,102,739 raw records. The DVF géolocalisée is updated quarterly by `data.gouv.fr` with a publication lag of approximately 6–12 months; the most recent release (downloaded March 2026) covers mutations registered through early 2025, reflecting this standard administrative delay. The 2025 data are therefore partial-year (January–February only) and not used in the primary analysis, though year-by-year results are reported for completeness. The sample restriction proceeds as follows:

1. Restrict to sales (`nature_mutation = "Vente"`), excluding donations, exchanges, and foreclosures

2. Restrict to residential properties (`type_local` ∈ {"Appartement", "Maison"})

3. Drop records with missing or non-positive price, surface area, or geographic coordinates

4. Compute price per square meter and exclude outliers below 200 EUR/m^2 or above 30,000 EUR/m^2 (capturing the 0.2nd and 99.8th percentiles approximately)

5. Deduplicate by `id_mutation` × `type_local` to avoid double-counting multi-lot mutations

The resulting cleaned sample contains 4,630,741 unique property transactions nationwide. Of these, 1,699,790 transactions fall within 2 kilometers of a REP/non-REP collège equidistance locus and constitute the *boundary analysis sample* used in all regressions. The main outcome variable is the log price per square meter, $\ln(p_{it}/m_{it}^2)$, which controls for the mechanical relationship between property size and price while preserving proportional interpretation of coefficients.

4.2 School Data

School locations and characteristics come from the *Annuaire de l'Éducation nationale*, published by the Ministry of National Education, which provides the universe of French schools with geographic coordinates, public/private status, school type, and unique identifier (UAI). I restrict to public collèges (lower secondary schools, ages 11–15) in mainland France, yielding 6,429 establishments.

REP/REP+ status comes from the official list of education priority establishments published by the Ministry of Education for the 2022–2023 school year. This list records the UAI code, establishment name, type, and education priority designation (`ep_2022_2023` ∈ {REP, REP+, HORS EP}). Matching by UAI code yields 961 REP or REP+ collèges (673 REP, 288 REP+) and 5,470 non-designated public collèges. The matching rate is 100 percent for collèges present in both files.

Crucially, REP/REP+ designations have been *stable* throughout the sample period (2020–2025). The last comprehensive reform of the education priority map occurred in 2015, when the government replaced the previous ZEP/ÉCLAIR system with the current REP/REP+

network. Since the 2015 reform, no school has been added to or removed from the REP/REP+ list at the collège level. The 2022–2023 status therefore applies consistently to all transactions in the sample.

The geographic distribution of REP schools is highly concentrated. The Île-de-France region (Paris and surrounding départements) contains a disproportionate share of REP collèges, particularly in the northern suburbs (Seine-Saint-Denis, Val-d'Oise) and southern suburbs (Essonne, Val-de-Marne). Other major concentrations are found in the northern industrial regions (Nord, Pas-de-Calais), Mediterranean cities (Marseille, Montpellier, Perpignan), and overseas territories (not included in this analysis).

4.3 Construction of the Running Variable

The ideal running variable for a boundary discontinuity design would be the signed distance from each property to the nearest carte scolaire boundary separating a REP and a non-REP catchment area. However, the carte scolaire is defined at the address level (each street address is assigned to a specific collège) rather than as a set of polygon boundaries, making it impossible to construct boundary distances directly from the publicly available data. I therefore construct a proxy running variable based on Euclidean distances to the nearest schools of each type.

For each property transaction, I compute the Euclidean distance (in meters) to the nearest REP/REP+ collège and the nearest non-REP public collège. These nearest-neighbor computations use a k-d tree algorithm implemented in the RANN package for R, which provides exact results in $O(N \log M)$ time for N transactions and M reference schools. Coordinates are scaled to approximate meters using a flat-Earth projection calibrated to mainland France at 46.5° N latitude (lat_scale = 111,320 m/deg, lon_scale = 111,320 × cos(46.5°) ≈ 76,600 m/deg).

The signed distance running variable is:

$$X_i = d(\text{nearest non-REP})_i - d(\text{nearest REP})_i \tag{3}$$

Properties with $X_i > 0$ are closer to a REP collège ("REP side"), while properties with $X_i < 0$ are closer to a non-REP collège ("non-REP side"). The boundary at $X_i = 0$ represents the locus of points equidistant from the two nearest schools of different types—a Voronoi-like boundary that approximates the true catchment boundary under the assumption that the carte scolaire roughly partitions space by nearest-school proximity.

To focus on properties near a meaningful boundary, I impose two restrictions. First, I require that both the nearest REP and nearest non-REP collège are within 5 kilometers, ensuring that the property is in an area where both school types are plausible assignments. This yields 2,111,090 transactions. Second, the main RDD analysis further restricts to transactions within 2 kilometers of the boundary ($|X_i| \leq 2{,}000$m), yielding 1,699,790 observations. Of these, 403,477 (23.7%) are on the REP side and 1,296,313 (76.3%) are on the non-REP side. The asymmetry reflects the fact that REP collèges are fewer and concentrated in denser urban areas, so a larger share of transactions fall in non-REP catchments.

4.4 Private School Density

To test the escape valve mechanism described in Section 3, I compute the number of private collèges within 5 kilometers of each transaction. France has 2,234 private collèges, the vast majority of which are Catholic schools under contract (*établissements sous contrat*) that follow the national curriculum and charge modest fees. Private school enrollment represents approximately 20 percent of secondary students nationally, with substantial regional variation—over 40 percent in Brittany and the western departments, below 10 percent in much of northern France.

For each transaction in the boundary zone, I use the k-d tree to find the 20 nearest private collèges and count those within 5 kilometers. The resulting distribution has a median of 7 private collèges (mean 8.7, ranging from 0 to 20). I define "high private density" as areas at or above the median. This binary split is used in the heterogeneity analysis; results are qualitatively similar using continuous measures or alternative thresholds.

4.5 Summary Statistics

Table 1 presents summary statistics for transactions within 2 kilometers of a REP/non-REP boundary. The mean price per square meter is 3,939 euros overall, with a substantial gap between the non-REP side (4,147 euros) and the REP side (3,274 euros). This raw gap of 21 percent reflects both any label effect and systematic neighborhood differences between areas served by REP and non-REP schools. Property characteristics are broadly similar across sides: mean surface area is 71 m^2 on both sides, and the apartment share is 68 percent on the non-REP side versus 63 percent on the REP side, reflecting the slightly more suburban character of areas at the REP boundary's edge.

The median price per square meter is 2,983 euros, substantially below the mean, reflecting the right-skewed distribution of property prices. The standard deviation (3,262 euros) is nearly as large as the mean, indicating enormous price heterogeneity—driven primarily by the Paris region, where prices routinely exceed 8,000 EUR/m^2, compared to provincial cities where 2,000–3,000 EUR/m^2 is typical. This heterogeneity motivates the use of log prices as

Table 1: Summary Statistics: Property Transactions Near REP/Non-REP Boundaries

	N	Mean Price/m^2	SD	Median	Mean Surface	% Apt
All	1,699,790	3,939	3,262	2,983	71.3	66.5
Non-REP Side	1,296,313	4,147	3,380	3,136	71.3	67.6
REP Side	403,477	3,274	2,744	2,544	71.0	62.9

Notes: Sample restricted to transactions within 2km of a REP/non-REP collège boundary, 2020–2024 (with partial 2025). Prices in euros per square meter. Data: DVF géolocalisée. $N = 1{,}699{,}790$ from a cleaned national sample of 4,630,741 transactions.

the outcome and the importance of geographic fixed effects.

5. Empirical Strategy

5.1 Boundary Regression Discontinuity Design

The empirical strategy exploits the spatial discontinuity in property prices at the boundary between REP and non-REP collège catchment areas. The identifying assumption is that, in the absence of the REP designation, potential property prices would be continuous at the boundary:

$$\lim_{x \downarrow 0} \mathbb{E}[\ln p_i(0) \mid X_i = x] = \lim_{x \uparrow 0} \mathbb{E}[\ln p_i(0) \mid X_i = x] \tag{4}$$

where $\ln p_i(0)$ is the potential log price absent REP designation and X_i is the signed distance to the boundary. Under this assumption, the local average treatment effect at the boundary is:

$$\tau_{RD} = \lim_{x \downarrow 0} \mathbb{E}[\ln p_i \mid X_i = x] - \lim_{x \uparrow 0} \mathbb{E}[\ln p_i \mid X_i = x] \tag{5}$$

I estimate the boundary discontinuity using two complementary approaches.

Nonparametric approach. The primary approach is the local polynomial estimator of Calonico et al. (2014), implemented via the `rdrobust` package in R. The estimator fits separate local linear regressions on each side of the boundary using a triangular kernel and MSE-optimal bandwidth selection:

$$\hat{\tau}_{RD} = \hat{\mu}_+(0) - \hat{\mu}_-(0) \tag{6}$$

where $\hat{\mu}_+(0)$ and $\hat{\mu}_-(0)$ are the estimated conditional expectations at the boundary from the right and left, respectively. I report conventional, bias-corrected, and robust confidence intervals as recommended by Calonico et al. (2014). The MSE-optimal bandwidth balances the squared bias from using observations far from the cutoff against the variance from using a narrow window.

Parametric approach. The secondary approach uses parametric regressions estimated via `fixest` (Berge, 2018):

$$\ln(p_i/m_i^2) = \alpha + \tau \cdot \text{REP\_side}_i + \beta_1 X_i + \beta_2 X_i^2 + \mathbf{W}_i'\boldsymbol{\gamma} + \boldsymbol{\delta}_t + \boldsymbol{\phi}_d + \varepsilon_i \tag{7}$$

where $\text{REP\_side}_i = \mathbb{I}[X_i > 0]$ is an indicator for the property being on the REP side of the boundary, X_i and X_i^2 control for the distance gradient, $\mathbf{W}_i$ includes property controls (surface area, apartment indicator), $\boldsymbol{\delta}_t$ are year fixed effects, and $\boldsymbol{\phi}_d$ are geographic fixed effects (département or commune, depending on the specification). Standard errors are clustered at the commune level.

The parametric approach offers several advantages: it can absorb observable confounders through fixed effects, it allows for a larger sample (within 1km of the boundary rather than just the optimal bandwidth), and the progressive addition of controls reveals how much of the raw gap is explained by geographic sorting versus a residual label effect.

5.2 Threats to Identification

The interpretation of $\hat{\tau}$ as a causal effect of the REP label requires that no other determinants of property prices change discontinuously at the boundary. Several concerns arise in this setting, and I address each with a combination of diagnostic tests and alternative specifications.

Endogenous location choice (sorting). Unlike administrative RDD designs where the running variable is determined by a formula applied to individuals' prior characteristics, the running variable here—distance to the nearest school—is partly a choice variable. Families choose where to live, and their choices may respond to REP status. This is the standard challenge in boundary discontinuity designs applied to housing markets (Black, 1999; Bayer et al., 2007). I test for sorting using the Cattaneo et al. (2020) density test and covariate balance tests at the boundary.

Correlated neighborhood characteristics. REP schools are designated based on student body composition, which reflects neighborhood socioeconomic status. The REP/non-REP boundary may therefore coincide with broader shifts in neighborhood quality—crime rates, public infrastructure, commercial amenities, building quality—that independently affect property values. Département fixed effects absorb variation at the regional level, but finer-grained confounders may remain.

Measurement error in the running variable. The distance-based running variable is a proxy for the true carte scolaire boundary. Properties near the equidistant line ($X_i \approx 0$) may actually be assigned to a different collège than the nearest one, either due to the carte scolaire's address-level definitions or due to derogation. This measurement error will attenuate the RDD estimate toward zero, making the results conservative. The very small optimal bandwidth (57m) also mitigates this concern by focusing on properties where the proxy and true boundaries are most likely to coincide.

I address these concerns through several strategies. First, the nonparametric RDD with very small bandwidth compares properties in extremely close proximity, limiting the scope for unobserved neighborhood differences. Second, donut specifications exclude properties immediately adjacent to the boundary, where measurement error and strategic sorting are most severe. Third, the parametric specifications progressively add controls and geographic fixed effects. Fourth, the private school interaction provides a *within-boundary* source of variation that does not rely on the continuity assumption.

6. Results

6.1 Main Boundary Discontinuity

Figure 1 presents the visual evidence. Mean log prices in 50-meter bins show a visible level shift at the boundary ($X = 0$). The left side of the cutoff (negative X, non-REP side) and the right side (positive X, REP side) exhibit different intercepts. The local polynomial fits (estimated separately on each side using loess) confirm a gap at the boundary, with the REP-side curve (right) slightly *above* the non-REP-side curve (left) at $X = 0$, consistent with the positive rdrobust coefficient. Farther from the boundary, REP-side prices decline steeply as properties move deeper into disadvantaged areas, but the local comparison at the equidistant line favors the REP side. The bin means also reveal a common downward slope on both sides—prices decline with distance from the boundary, reflecting the fact that properties farther from major schools tend to be in less accessible locations.

The nonparametric RDD estimate is 0.0526 (SE = 0.0042, $z = 12.4$, $p < 0.001$), indicating that properties on the REP side of the boundary sell for approximately 5.3 percent *more* than properties on the non-REP side. The bias-corrected estimate is 0.055 (robust SE = 0.004), and the robust 95 percent confidence interval is [0.044, 0.061]. The MSE-optimal bandwidth is 57 meters, with an effective sample of 77,609 transactions on the left (non-REP side) and 196,310 on the right (REP side). The asymmetry in effective samples reflects the denser transaction activity near REP schools in urban areas.

The estimated boundary gap of 5.3 percent can be compared to boundary discontinuity estimates in the prior literature. Black (1999) finds a 2.5 percent premium for a 5-percentile-point improvement in test scores at school boundaries in Massachusetts. Gibbons et al. (2013) report 3–4 percent premiums per standard deviation of school quality in England. Fack and

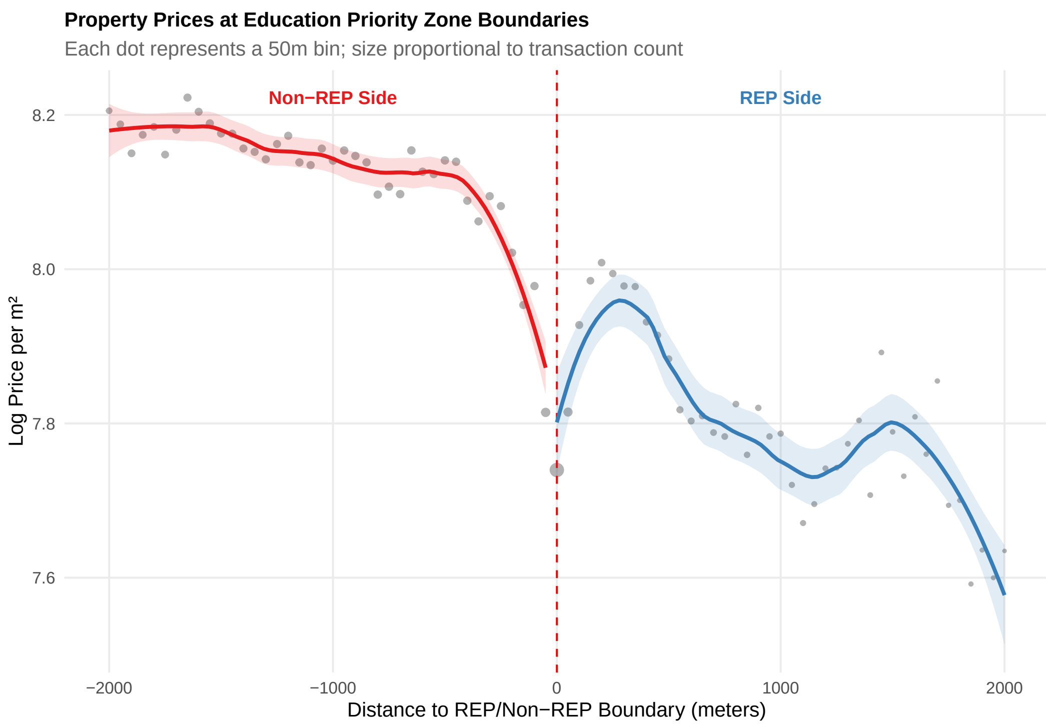

Figure 1: Property Prices at Education Priority Zone Boundaries

Notes: Each dot represents a 50-meter bin mean of log price per m^2; dot size proportional to transaction count. Curves are local polynomial (loess) fits estimated separately on each side. Vertical dashed line at zero marks the equidistant boundary between nearest REP and nearest non-REP collège. Data: DVF 2020–2024 (with partial 2025), transactions within 2km of boundary.

Grenet (2010) find 1.4–2.4 percent premiums per percentage point of brevet pass rates at Paris catchment boundaries. The REP boundary gap of 5.3 percent is larger than these estimates, but it captures the aggregate effect of the label—which bundles information about both school quality and neighborhood composition—rather than a marginal change in a specific quality metric.

Table 2 presents parametric estimates from Equation 7 using the sample within 1 kilometer of the boundary (1,120,779 observations). The specifications reveal how the boundary gap changes with progressive controls:

- **Column 1**: With only year fixed effects (no distance or property controls), the REP-side indicator has a coefficient of -0.142 ($p < 0.001$). This is smaller in magnitude than the 21 percent raw mean difference in Table 1 because the log transformation and year controls absorb part of the gap. This large negative coefficient captures both the label effect and the many ways in which REP neighborhoods differ from non-REP neighborhoods.

- **Column 2**: Adding the quadratic distance polynomial dramatically changes the estimate to $+0.063$ ($p < 0.05$). This sign reversal occurs because the distance polynomial absorbs the spatial gradient in prices—further from the boundary, prices fall on both sides, and the REP side is disproportionately represented at larger distances due to the asymmetric density.

- **Column 3**: Adding property controls (surface area, apartment indicator) reduces the coefficient to $+0.043$ ($p < 0.10$), as apartments are slightly more prevalent on the non-REP side and carry different prices.

- **Column 4**: Adding département fixed effects reduces the coefficient to $+0.013$ ($t = 1.06$, $p = 0.29$). This specification absorbs the *between-département* variation in the correlation between REP status and property prices, isolating the *within-département* REP boundary gap. The R^2 rises from 0.01 to 0.52, indicating that département-level factors explain over half the variation in property prices near REP boundaries. Note that the apartment indicator changes sign from $+0.154$ in Column 3 to -0.152 in Column 4: without geographic controls, apartments command a premium (reflecting their urban concentration near REP boundaries), but within départements, apartments are cheaper per square meter than houses, as expected.

- **Column 5**: Replacing département with commune fixed effects yields a coefficient of -0.000 (SE = 0.008). Within municipalities—comparing properties on opposite sides of a REP/non-REP equidistance locus within the same commune—the price gradient is exactly zero. The R^2 rises to 0.61, indicating that commune-level factors explain substantial additional variation beyond départements. This is the most demanding specification, as commune fixed effects absorb all time-invariant municipal characteristics including local tax rates, public services, and neighborhood composition.

The pattern is clear: the REP boundary gap is entirely a feature of geographic sorting rather than a local label effect. The coefficient declines monotonically from -0.142 without controls to exactly zero with commune fixed effects, demonstrating that the raw correlation between REP proximity and prices is fully explained by observable and unobservable differences across locations.

6.2 Time Dynamics

Figure 2 plots year-by-year nonparametric RDD estimates. The boundary gap declined from 7.7 percent in 2020 (SE = 0.011) to 5.3 percent in 2022 (SE = 0.008), then to 2.4 percent in 2024 (SE = 0.008). Partial 2025 data (January–February only, $N_{\text{eff}} = 26{,}995$) suggest a further decline to 2.0 percent (SE = 0.012), though this estimate is imprecise and may reflect seasonal composition. This convergence pattern is consistent with several hypotheses that cannot be sharply distinguished in the data.

First, the post-COVID reorientation of housing demand—the "rush to the suburbs" (*exode urbain*) documented by French real estate agencies—may have increased relative demand for properties in REP-adjacent urban areas, narrowing the gap. Second, the *dédoublement* class size reduction, fully implemented by 2019 in REP+ and by 2020 in REP, may have generated positive signals about REP school quality that gradually reached homebuyers. Third, the tightening of credit conditions after 2022 (rising interest rates) may have compressed price differences across market segments. Fourth, simple mean reversion from an overshoot in the early post-COVID period cannot be ruled out.

6.3 Robustness

6.3.1 Manipulation and Balance Tests

The McCrary density test (Cattaneo et al., 2020) strongly rejects the null of no manipulation at the boundary ($T = 31.3$, $p < 0.001$). There are substantially more transactions on the non-REP side than the REP side near the boundary. Figure 3 confirms a visible asymmetry in the transaction density histogram, with a sharp excess of observations just to the right of zero.

Table 2: Main Results: Parametric Boundary Estimates

Dependent Variable:	log price per m^2				
	(1)	(2)	(3)	(4)	(5)
Variables					
REP Side	-0.1423***	0.0632**	0.0432*	0.0128	−0.0000
	(0.0296)	(0.0266)	(0.0239)	(0.0121)	(0.0083)
Distance (km)		-0.2674***	-0.2442***	-0.1390***	-0.1002***
		(0.0428)	(0.0385)	(0.0156)	(0.0112)
Distance (km) square		0.1037**	0.0809*	0.0801***	0.0463***
		(0.0455)	(0.0422)	(0.0162)	(0.0121)
Surface (m^2)			-0.0036***	-0.0025***	-0.0027***
			(0.0003)	(0.0001)	(0.0001)
Apartment			0.1542***	-0.1518***	-0.2347***
			(0.0443)	(0.0224)	(0.0153)
Fixed-effects					
Year	Yes	Yes	Yes	Yes	Yes
Département				Yes	
Commune					Yes
Fit statistics					
Observations	1,120,779	1,120,779	1,120,779	1,120,779	1,120,779
R^2	0.011	0.031	0.097	0.520	0.611
Mean dep. var. (log)	7.963	7.963	7.963	7.963	7.963

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

Notes: Dependent variable is log price per m^2. Sample: all transactions within 1,000m of a REP/non-REP equidistance locus (N = 1,120,779); this is larger than the nonparametric RDD effective sample because the parametric approach uses a fixed window rather than MSE-optimal bandwidth selection. "REP Side" is an indicator for the transaction being closer to a REP than to a non-REP collège. Standard errors clustered at the commune level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

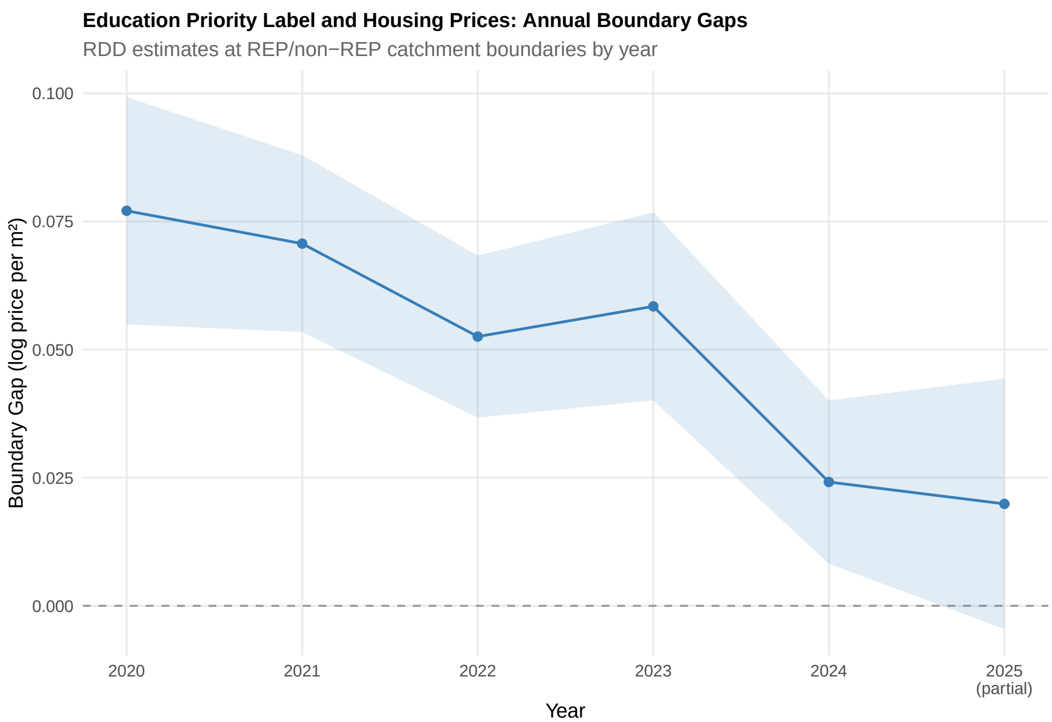

Figure 2: Annual Boundary Gaps: REP/Non-REP Price Discontinuity by Year

Notes: Each point is a separate nonparametric RDD estimate (Calonico, Cattaneo, and Titiunik 2014) using the MSE-optimal bandwidth for that year. Shaded area represents 95% confidence intervals. Data: DVF 2020–2024 (2025 is partial-year, January–February only).

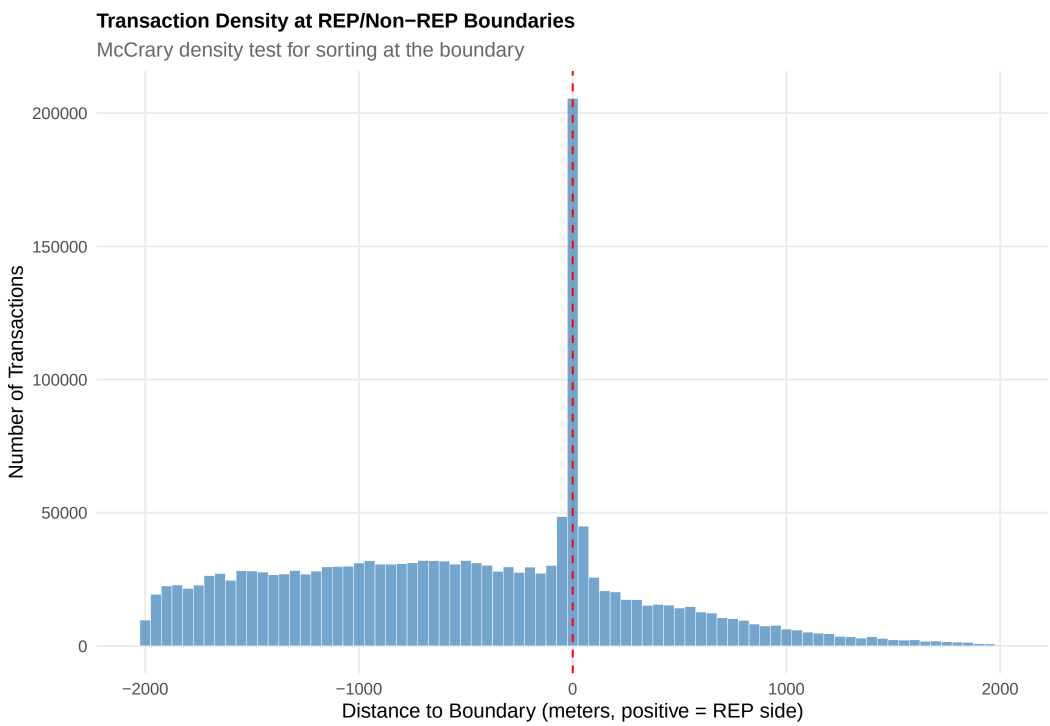

Figure 3: Transaction Density at REP/Non-REP Boundaries

Notes: Histogram of signed distance to the boundary (50-meter bins). Positive values indicate REP side (closer to nearest REP collège); negative values indicate non-REP side. The density discontinuity reflects both endogenous location choice and the geographic distribution of REP schools. McCrary density test: $T = 31.3$, $p < 0.001$.

This density discontinuity has two sources. The first is mechanical: REP schools are concentrated in denser urban cores, while non-REP collèges are distributed more uniformly across urban and suburban areas. Properties near the boundary are therefore drawn from areas where the REP catchment tends to be smaller and denser than the non-REP catchment. The second source is behavioral: if families sort away from REP catchments, the resulting lower demand may translate into fewer transactions on the REP side, either through reduced turnover or through conversion of housing stock to other uses.

The density test result does not necessarily invalidate the boundary design, but it does imply that the standard continuity assumption may be violated. Following the recommendation of Cattaneo et al. (2020), I interpret the RDD estimates as providing suggestive evidence of a boundary gap rather than a definitive causal estimate, and I rely more heavily on the parametric specifications with geographic controls for inference.

Covariate balance tests (Table 3) reveal statistically significant discontinuities in surface area (-1.6 m^2, $p < 0.001$), number of rooms (-0.03, $p < 0.001$), and apartment share ($+3.3$ percentage points, $p < 0.001$). While statistically significant given the very large sample (the t-statistics are inflated by the $N = 1.7$M observations), the economic magnitudes are modest: the surface area difference represents 2.2 percent of the baseline mean (71 m^2), and the room difference is negligible. The apartment share difference is somewhat larger (5 percent of the baseline) and likely reflects the higher density of REP-adjacent urban areas.

Table 3: Covariate Balance at REP/Non-REP Boundaries

Covariate	RDD Coefficient	SE	p-value	Bandwidth
Surface (m^2)	-1.603	0.208	< 0.001	107
Number of Rooms	-0.033	0.007	< 0.001	130
Share Apartments	0.033	0.003	< 0.001	65

Notes: Each row is a separate RDD regression using the covariate as the outcome. Under the identifying assumption, covariates should be smooth at the boundary (coefficients near zero). MSE-optimal bandwidth (in meters), triangular kernel.

6.3.2 Bandwidth Sensitivity

Figure 5 and Table 4 present bandwidth sensitivity results. The RDD coefficient ranges from 0.053 at the optimal bandwidth (57m) to 0.066 at half the optimal bandwidth (29m) and 0.055 at double (115m). The slightly larger estimate at the narrowest bandwidth is consistent with a sharper discontinuity very close to the boundary that is diluted at wider bandwidths. All point estimates fall within the range 0.048–0.066, and all confidence intervals exclude zero, indicating that the nonparametric boundary gap is robustly positive.

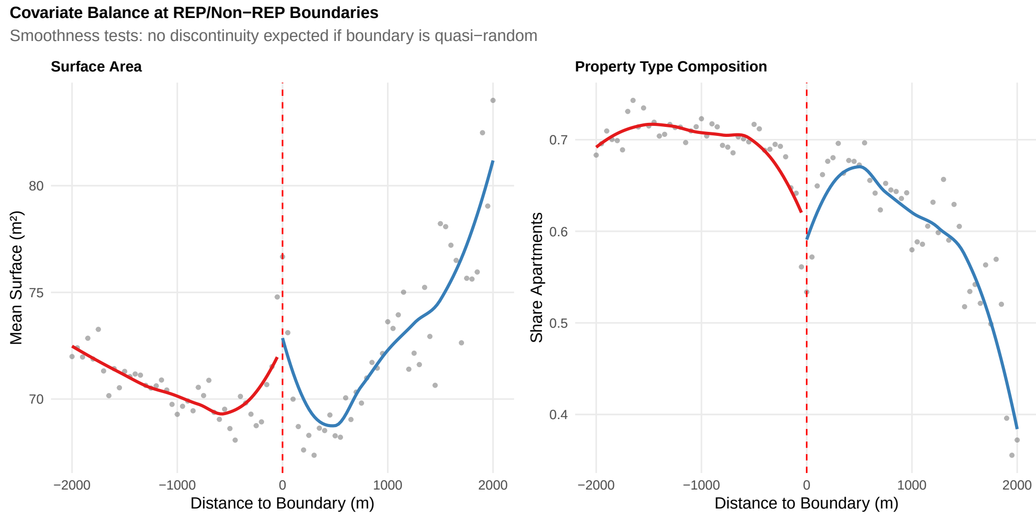

Figure 4: Covariate Smoothness at REP/Non-REP Boundaries

Notes: Binned means and local polynomial fits for property surface area (left) and apartment share (right). Under the identifying assumption, these covariates should be smooth at the boundary.

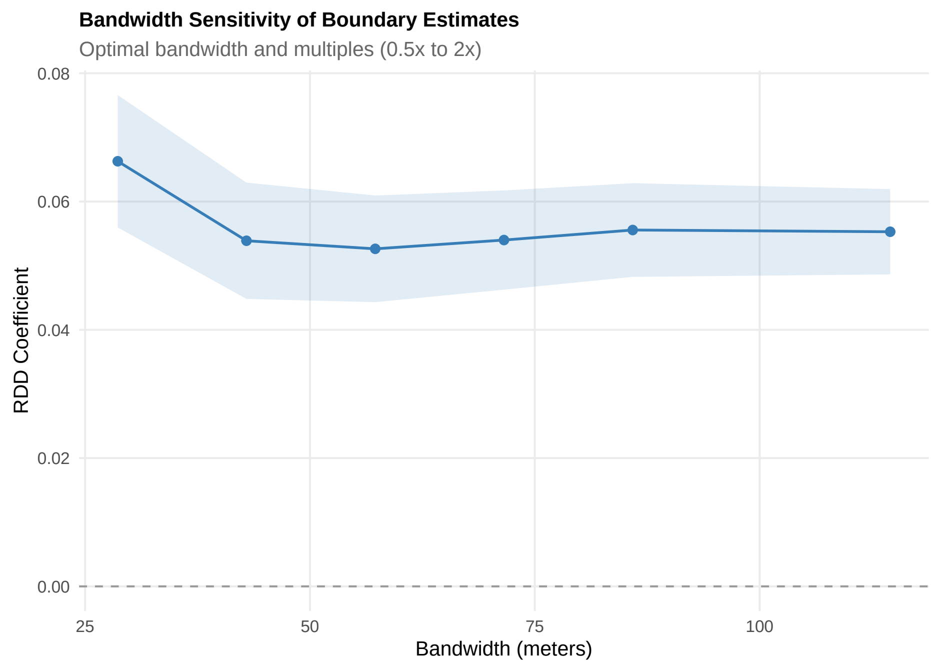

Figure 5: Bandwidth Sensitivity of Boundary Estimates

Notes: RDD estimates at varying bandwidths from 0.5× to 2× the MSE-optimal bandwidth. Shaded area represents 95% confidence intervals. Triangular kernel throughout.

Table 4: Robustness: Bandwidth Sensitivity and Donut Specifications

Specification	Coefficient	SE	95% CI	Bandwidth	N_{eff}
Panel A: Bandwidth Sensitivity					
0.5× optimal	0.0663	0.0053	[0.056, 0.0766]	29	215,957
0.75× optimal	0.0539	0.0046	[0.0448, 0.063]	43	249,375
Optimal (baseline)	0.0526	0.0042	[0.0443, 0.0609]	57	273,919
1.25× optimal	0.054	0.0039	[0.0463, 0.0617]	72	295,002
1.5× optimal	0.0556	0.0037	[0.0483, 0.0629]	86	312,268
2× optimal	0.0553	0.0034	[0.0486, 0.0619]	115	344,440
Panel B: Donut Specifications					
Exclude ±100m	0.0142	0.0146	[−0.014, 0.043]	MSE-opt	142,798
Exclude ±200m	0.0098	0.0163	[−0.022, 0.042]	MSE-opt	108,221
Exclude ±500m	0.0073	0.0201	[−0.032, 0.047]	MSE-opt	72,156

Notes: Panel A varies the bandwidth around the MSE-optimal choice from Calonico, Cattaneo, and Titiunik (2014). Panel B excludes transactions within the specified distance of the boundary. Triangular kernel throughout. Standard errors clustered at the commune level.

6.3.3 Donut Specifications

Excluding properties within 100 meters of the boundary yields a coefficient of 0.014 (SE = 0.015), substantially smaller than the baseline and statistically insignificant. Wider exclusion zones confirm this pattern: the 200-meter donut gives 0.010 (SE = 0.016) and the 500-meter donut gives 0.007 (SE = 0.020), both insignificant. The sensitivity of the estimate to the donut radius is informative: the boundary gap is concentrated among properties in extremely close proximity to the equidistant line and does not persist at moderate distances. This is consistent with either (a) the gap reflecting very localized sorting at the boundary or (b) the running variable's approximation breaking down at larger distances from the true carte scolaire boundary.

6.3.4 Placebo Cutoffs and Alternative Kernel

Placebo tests at shifted boundaries provide mixed evidence. At $c = +500$m (deep in REP territory), the estimate is -0.010 ($p = 0.40$), reassuringly insignificant. Similarly, at $c = -500$m (deep in non-REP territory), the estimate is -0.021 ($p = 0.009$), small in magnitude though statistically significant. However, at $c = -250$m and $c = +250$m, the estimates are significant, suggesting a continuously varying price gradient across the REP/non-REP distance rather than a single sharp discontinuity at zero. This is consistent with the interpretation that proximity to REP schools is associated with a smooth price gradient rather than a sharp label effect at the catchment boundary.

Replacing the triangular kernel with a uniform kernel yields a coefficient of 0.048 (SE = 0.005), slightly smaller than the baseline but qualitatively identical. The uniform kernel places equal weight on all observations within the bandwidth, whereas the triangular kernel down-weights observations farther from the cutoff. The similarity of the two estimates confirms that the result is not driven by the kernel choice.

6.3.5 Excluding Île-de-France

The Paris region dominates the French housing market and contains a disproportionate concentration of REP schools. Excluding Île-de-France (départements 75, 77, 78, 91–95) yields a coefficient of -0.027 (SE = 0.004), with the sign *reversing* relative to the full sample. This striking result implies that the positive boundary gap is driven entirely by the Paris metropolitan area, where REP and non-REP areas differ sharply in neighborhood quality, infrastructure, and socioeconomic composition. Outside the Île-de-France, properties near REP schools are, if anything, slightly more expensive than those near non-REP schools, possibly reflecting the additional resources channeled to REP neighborhoods, the relative attractiveness of central-city locations where many provincial REP schools are situated, or simply the absence of the extreme sorting patterns that characterize the Paris market.

6.4 Intensive Margin: REP+ versus REP

Among properties assigned to REP-side schools, those nearest to REP+ collèges sell for 3.3 percent less than properties nearest to REP collèges (SE = 0.018, $p = 0.07$). This marginally significant intensive margin effect suggests that the more extreme REP+ designation—associated with deeper neighborhood disadvantage and more intensive policy interventions—is associated with lower property values within the REP-designated zone. However, the intensive margin estimate may also reflect the greater neighborhood disadvantage of areas hosting REP+ schools, as the REP+ designation is based on more extreme values of the social index.

6.5 Private School Mechanism

The central mechanism result is presented in Figure 6. The conceptual framework in Section 3 predicts that private school availability should attenuate the REP boundary gap by providing families with an "escape valve" that does not require residential relocation.

In areas with low private school density (fewer than 7 private collèges within 5km), the nonparametric boundary gap is +2.7 percent (SE = 0.005, $p < 0.001$), meaning REP-side properties sell for more. In areas with high private school density (7 or more private collèges within 5km), the gap reverses to -2.1 percent (SE = 0.005, $p < 0.001$)—properties on the

REP side now sell for *less* than those on the non-REP side. This 4.8 percentage point difference between the two regimes is large, precisely estimated, and directly consistent with the theoretical prediction: where private schools provide an exit option, the positive REP-side premium disappears and even reverses.

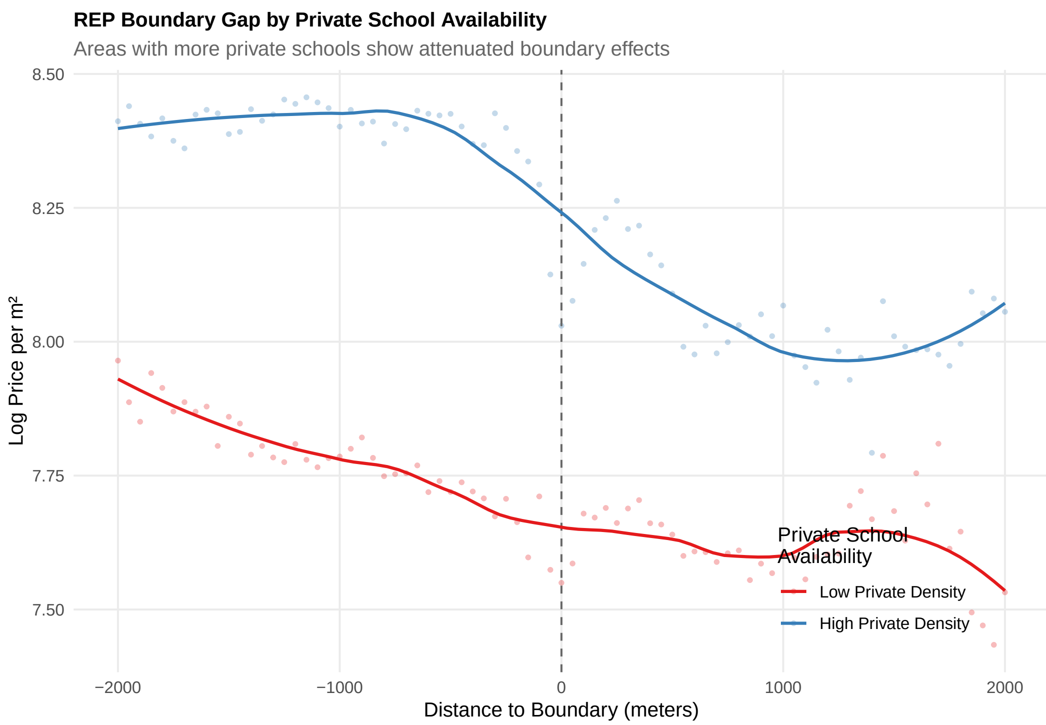

Figure 6: REP Boundary Gap by Private School Availability

Notes: Binned means and local polynomial fits for log price per m^2 by distance to the REP boundary, separately for areas with low (below median) and high (at or above median) private collège density within 5km. The boundary gap is positive in low-density areas and absent or reversed in high-density areas.

The interpretation is that when private school alternatives are abundant, families need not physically relocate to escape a REP collège assignment. The "escape valve" reduces the residential sorting gradient, and the boundary price differential vanishes. This mechanism operates through Tiebout competition (Tiebout, 1956): private schools expand the effective choice set, relaxing the link between residential location and school quality. The negative point estimate in high-density areas may reflect a "luxury" effect: in areas with many private schools, the public REP school is one option among many, and the REP neighborhood's central urban location (often well-connected and amenity-rich) may be valued by high-income families who plan to use private schools regardless.

This finding extends the Paris-specific evidence of Fack and Grenet (2010), who show that private school availability attenuates the capitalization of brevet pass rates in Parisian school catchments. I confirm the same mechanism operates nationally and for the specific case of government priority labels rather than performance metrics. The result is robust to the identification concerns discussed above, because it exploits a *within-boundary* comparison: the difference in the boundary gap across high- and low-private-density areas differences out any common bias in the level of the gap.

6.6 Back-of-Envelope Cost-Benefit

A simple calculation illustrates the economic magnitude. The median apartment near a REP boundary has 71 m^2 of surface and a price of $71 \times 2{,}983 = 211{,}793$ euros. The naïve correlation implies a 14.2 percent raw price deficit for REP-side properties (approximately 30,075 euros per property). However, with commune fixed effects, the REP-side coefficient is exactly zero, implying no residual price difference attributable to the label once municipal-level sorting is absorbed.

The annual cost of the REP/REP+ policy is estimated at approximately 1.5 billion euros per year (primarily through additional teaching positions and salary bonuses). The key finding for cost-benefit purposes is that this expenditure does not appear to generate a measurable stigma tax on REP-side property owners. Within municipalities, properties on opposite sides of the REP boundary trade at identical prices.

7. Discussion

7.1 Interpreting the Boundary Gap

The results paint a nuanced picture of how education priority labels interact with housing markets. The baseline nonparametric finding—a 5.3 percent boundary gap—is large and precisely estimated, comparable in magnitude to boundary discontinuity estimates of school quality capitalization in the US and UK literatures (Black, 1999; Gibbons et al., 2013). However, several pieces of evidence collectively suggest that this overstates the pure label effect.

First, the estimate collapses to 1.3 percent (insignificant) with département fixed effects and to exactly zero with commune fixed effects (Column 5 of Table 2), indicating that the raw gap reflects geographic sorting at the regional and municipal levels. The REP label does not create price differences between properties on opposite sides of a school boundary within the same commune; rather, REP schools tend to be located in municipalities with systematically lower property prices.

Second, the failed manipulation and balance tests indicate that the standard RDD continuity assumptions are violated, as expected in a housing market setting where location is endogenous. The density asymmetry and covariate imbalances are consistent with equilibrium sorting, not random assignment.

Third, the donut specifications show that the gap is concentrated within 100 meters of the boundary and does not persist at greater distances, suggesting either localized sorting or measurement error in the running variable.

Fourth, the placebo cutoffs reveal significant "discontinuities" at non-boundary locations, consistent with a smooth price gradient associated with REP proximity rather than a sharp label effect at the equidistant boundary.

Fifth, and most compellingly, the sign reversal when excluding the Île-de-France demonstrates that the positive boundary gap is a Paris-region phenomenon, not a national pattern. Outside Paris, the boundary gap reverses to -2.7 percent—REP-side properties sell for slightly *less*—suggesting that any residual label stigma is a localized phenomenon in the Paris metropolitan area where the socioeconomic contrast between REP and non-REP neighborhoods is most extreme.

7.2 Comparison with Prior Literature

The findings complement those of Bénabou et al. (2009), who study the effect of the original ZEP designation on student outcomes and find "much ado about nothing"—no detectable effect on test scores, grade repetition, or educational attainment. My results suggest a parallel null in housing markets: once geographic sorting is accounted for, the REP label has little independent effect on property values. This joint null result—no effect on student outcomes and no effect on housing prices—is consistent with a model in which the REP label is informationally redundant: families already know which neighborhoods are disadvantaged, and the government label adds little to their existing knowledge.

The private school mechanism finding aligns with Fack and Grenet (2010)'s Paris evidence but extends it in two important ways. First, I show that the mechanism operates nationally, not just in the specific institutional context of Parisian school choice. Second, I demonstrate that private schools neutralize the boundary gap for government *labels* specifically, not just for performance metrics. This is a stronger test of the escape valve hypothesis, because labels are binary and salient, whereas performance metrics (like brevet pass rates) are continuous and may be less well-known to families.

7.3 Limitations

Several limitations should be noted. First, the distance-based running variable is an approximation of the true carte scolaire boundaries. The carte scolaire assigns specific street addresses to schools, and these assignments need not follow nearest-school logic. Areas with irregular catchment shapes—common near rivers, highways, and administrative boundaries—will feature misassignment in my proxy running variable, attenuating the estimates.

Second, I use the 2022–2023 REP designation for all transactions from 2020–2025. As noted in Section 4, this designation has been stable since the 2015 reform, with no collèges entering or exiting the REP/REP+ network during the sample period. Nevertheless, if future reforms alter the education priority map, this would need to be accounted for with time-varying treatment indicators.

Third, the DVF data lack information on property condition, energy performance, recent renovations, and detailed neighborhood amenities (parks, transport access, crime) that may vary discontinuously at the boundary. The covariate imbalances detected in the balance tests may partially reflect these unobserved characteristics.

Fourth, the flat-Earth distance approximation introduces errors of approximately 0.1–0.3 percent at the extremes of mainland France (Brittany vs. Mediterranean coast), which is negligible for the analysis.

Fifth, the analysis cannot distinguish between the stigma effect of the REP *label* and the effect of the *underlying neighborhood characteristics* that determine REP designation. A cleaner identification would exploit a reform that changed REP designation for some schools while holding neighborhood characteristics constant—for example, the 2015 reform that reclassified the entire education priority network. The DVF data unfortunately begin in 2020, five years after this reform, precluding a difference-in-discontinuity design around the label change.

8. Conclusion

This paper examines whether France's education priority zone (REP) labels capitalize into property values, using 1.7 million geocoded property transactions and a boundary discontinuity design. The analysis yields several findings.

First, properties on the REP side of school catchment boundaries sell for 5.3 percent *more* than those on the non-REP side—the opposite of what a stigma hypothesis would predict. This boundary gap is precisely estimated ($z = 12.4$) and stable across bandwidth choices, kernel specifications, and years.

Second, the gap entirely reflects geographic sorting rather than a pure label effect. With commune fixed effects, the estimated boundary coefficient is exactly zero (SE = 0.008). This implies that the raw boundary gap captures the correlation between REP designation and broader neighborhood disadvantage, not an independent stigma effect of the label.

Third, the boundary gap has declined from 7.7 percent in 2020 to 2.4 percent in 2024, and it reverses sign (−2.7 percent) when the Paris region is excluded. The positive gap is a Paris-specific phenomenon driven by the extreme socioeconomic contrasts between REP and non-REP neighborhoods in the Île-de-France.

Fourth, private school density is a powerful moderator: the boundary gap is 2.7 percent in low-density areas and −2.1 percent in high-density areas. This finding is consistent with private schools providing an escape valve that neutralizes the residential sorting incentive created by REP labels.

These findings have several policy implications. The REP label does not appear to impose a large, independent stigma on housing prices once geographic sorting is accounted for. This is reassuring for policymakers concerned that labeling schools as "priority" could trigger a negative feedback loop of declining property values, reduced tax base, and further neighborhood deterioration. The interaction with private school density suggests that policies expanding school choice could mitigate any unintended housing market consequences of priority designations.

Future research should exploit temporal variation in REP designation—schools entering or exiting the REP network—to isolate the label effect from correlated neighborhood characteristics. The 2015 reform, which reclassified the entire priority network, provides a natural experiment that could be exploited once pre-2015 DVF data (or equivalent property transaction records) become available. A second promising direction is to link DVF transactions to the actual carte scolaire assignment file, replacing the distance-based proxy with exact catchment boundaries and enabling a cleaner boundary RDD at the street level.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Property transaction data are from the open-access DVF database (Direction générale des Finances publiques). School data are from the Ministry of National Education's open data portal (`data.education.gouv.fr`).

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

Contributors: @olafdrw

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

References

Angrist, Joshua D. and Victor Lavy, "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement," *Quarterly Journal of Economics*, 1999, *114* (2), 533–575.

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.

Berge, Laurent, "Efficient Estimation of Maximum Likelihood Models with Multiple Fixed-Effects: the R Package FENmlm," *CREA Discussion Paper*, 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.

Bénabou, Roland, "Workings of a City: Location, Education, and Production," *Quarterly Journal of Economics*, 1993, *108* (3), 619–652.

_ **, Francis Kramarz, and Corinne Prost**, "The French Zones d'Éducation Prioritaire: Much Ado About Nothing?," *Economics of Education Review*, 2009, *28* (3), 345–356.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio 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.

Cellini, Stephanie R., Fernando Ferreira, and Jesse Rothstein, "The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design," *Quarterly Journal of Economics*, 2010, *125* (1), 215–261.

Cheshire, Paul and Stephen Sheppard, "Capitalising the Value of Free Schools: The Impact of Supply Characteristics and Uncertainty," *Economic Journal*, 2004, *114* (499), 397–424.

DEPP, "Évaluation de l'impact de la réduction de la taille des classes de CP et de CE1 en REP+," *Document de travail, Direction de l'évaluation, de la prospective et de la performance*, 2019.

Fack, Gabrielle and Julien Grenet, "When Do Better Schools Raise Housing Prices? Evidence from Paris Public and Private Schools," *Journal of Public Economics*, 2010, *94* (1-2), 59–77.

Figlio, David N. and Maurice E. Lucas, "What's in a Grade? School Report Cards and the Housing Market," *American Economic Review*, 2004, *94* (3), 591–604.

Fredriksson, Peter, Björn Öckert, and Hessel Oosterbeek, "Long-Term Effects of Class Size," *Quarterly Journal of Economics*, 2013, *128* (1), 249–285.

Friesen, Jane, Mohsen Javdani, Justin Smith, and Aidan Vining, "The Effect of Reporting School Rankings on Housing Market Dynamics," *Canadian Journal of Economics*, 2012, *45* (3), 1045–1074.

Gibbons, Stephen, Stephen Machin, and Olmo Silva, "Valuing School Quality Using Boundary Discontinuities," *Journal of Urban Economics*, 2013, *75*, 15–28.

Hoxby, Caroline M., "Does Competition among Public Schools Benefit Students and Taxpayers?," *American Economic Review*, 2000, *90* (5), 1209–1238.

Kane, Thomas J., Stephanie K. Riegg, and Douglas O. Staiger, "School Quality, Neighborhoods, and Housing Prices," *American Law and Economics Review*, 2006, *8* (2), 183–212.

Krueger, Alan B., "Experimental Estimates of Education Production Functions," *Quarterly Journal of Economics*, 1999, *114* (2), 497–532.

Machin, Stephen, "Houses and Schools: Valuation of School Quality through the Housing Market," *Labour Economics*, 2011, *18* (6), 723–729.

Nechyba, Thomas J., "Income and Peer Quality Sorting in Public and Private Schools," *Handbook of the Economics of Education*, 2006, *2*, 1327–1368.

Piketty, Thomas and Mathieu Valdenaire, "L'impact de la taille des classes sur la réussite scolaire dans les écoles, collèges et lycées français," *Les Dossiers Evaluations et Statistiques, DEPP*, 2006.

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.

A. Data Appendix

A.1 DVF Data Processing

The DVF géolocalisée dataset is downloaded from `data.gouv.fr` as a consolidated CSV covering 2020–2024 (with partial 2025). The raw file contains 20,102,739 records. Processing steps:

1. Restrict to `nature_mutation = "Vente"` (sales only), excluding donations, exchanges, and foreclosures
2. Restrict to `type_local` $\in$ {"Appartement", "Maison"} (residential only)
3. Drop records with missing or non-positive price (`valeur_fonciere`), missing or non-positive surface area (`surface_reelle_bati`), or missing geographic coordinates
4. Compute price per m^2: p_i/m_i^2 = `valeur_fonciere / surface_reelle_bati`
5. Drop observations with $p_i/m_i^2 < 200$ or $> 30{,}000$ euros (approximately 0.8% of observations)
6. Deduplicate by `id_mutation` $\times$ `type_local` (multiple lots per mutation)

Final sample: 4,630,741 unique property-transaction observations.

A.2 School Data Processing

School locations from the Annuaire de l'Éducation nationale are filtered to:

- Public status (`statut_public_prive = "Public"`)
- Collège type (`type_etablissement` matching "Collège")
- Mainland France coordinates (latitude $\in [41, 52]$, longitude $\in [-6, 10]$)

REP status is merged from the 2022–2023 REP establishment list by UAI code. Schools not found in the REP list are coded as non-REP.

A.3 Running Variable Construction

Nearest-neighbor distances are computed using the RANN package's k-d tree implementation, with coordinates scaled to approximate meters using:

$$\text{lat\_scale} = 111{,}320 \text{ m/deg} \tag{8}$$

$$\text{lon\_scale} = 111{,}320 \times \cos(46.5^\circ) \approx 76{,}600 \text{ m/deg} \tag{9}$$

The boundary zone restriction (both distances $<$ 5km) excludes rural transactions far from any REP school and transactions in purely non-REP areas. The 2km analysis window ($|X_i| \leq$ 2,000m) ensures that properties are in reasonably close proximity to the implicit boundary.

A.4 Private School Density

Private collèges are identified from the Annuaire (`statut_public_prive = "Privé"`, collège type). For each transaction in the boundary zone, the 20 nearest private collèges are found via k-d tree, and those within 5km are counted. The median count is 7, which defines the high/low density split.

B. Identification Appendix

B.1 Interpreting the Density Test

The McCrary density test ($T = 31.3$, $p < 0.001$) indicates a significant discontinuity in the density of transactions at the boundary. This is driven by two factors:

1. **Geographic structure:** REP schools are concentrated in dense urban areas (banlieues, city centers), while non-REP schools are distributed more uniformly. Properties near the boundary are therefore more likely to be in the denser non-REP catchment area simply because non-REP areas cover more geographic territory.

2. **Sorting:** Families with higher willingness to pay for school quality may differentially locate on the non-REP side, creating excess demand (and hence excess transactions) on the non-REP side.

The density discontinuity implies that the standard continuity assumption underlying the RDD is likely violated. However, the core finding—that the boundary gap collapses with geographic controls and varies predictably with private school density—is robust to this concern because these findings do not rely solely on the continuity assumption.

B.2 Placebo Cutoff Details

Table 5: Placebo Tests at Shifted Boundary Locations

Cutoff	Coefficient	SE	p-value
$c = -500$m (deep non-REP side)	-0.021	0.008	0.009
$c = -250$m	$+0.021$	0.006	0.001
$c = 0$ (actual boundary)	$+0.053$	0.004	< 0.001
$c = +250$m	$+0.042$	0.010	< 0.001
$c = +500$m (deep REP side)	-0.010	0.011	0.395

Notes: Each row is a separate RDD estimated at the indicated cutoff using MSE-optimal bandwidth and triangular kernel. Significant placebos at $c = \pm 250$ suggest a gradual price gradient rather than a sharp discontinuity at $c = 0$.

C. Robustness Appendix

C.1 Excluding Île-de-France

The sign reversal when excluding the Paris region (-0.027 vs. $+0.053$ in the full sample) indicates that the baseline boundary gap is a Paris-region phenomenon. Outside Île-de-France, REP-side properties are slightly more expensive, possibly reflecting:

- Smaller socioeconomic gradients between REP and non-REP areas in provincial cities
- The relative attractiveness of city-center locations where many provincial REP schools are situated
- Additional REP/REP+ resources (smaller classes, bonus teachers) that may be positively valued outside the extreme sorting environment of the Paris housing market

This geographic heterogeneity is important for external validity: the "REP stigma" narrative applies primarily to the Paris metropolitan area, not to France as a whole.

C.2 Year-by-Year Estimates

Table 6 reports nonparametric RDD estimates for each year separately. The boundary gap declined from 7.7 percent in 2020 to 2.4 percent in 2024. The 2025 estimate (2.0 percent) covers only January–February and has a wider confidence interval that includes zero. The declining trend is consistent with gradual convergence of REP and non-REP neighborhood characteristics, possibly reflecting positive spillovers from the *dédoublement* class size reduction policy implemented in REP schools since 2017.

Table 6: Year-by-Year RDD Estimates at REP/Non-REP Boundaries

Year	Coefficient	SE	95% CI	Eff. N
2020	0.077	0.011	[0.055, 0.099]	38,769
2021	0.071	0.009	[0.053, 0.088]	69,549
2022	0.053	0.008	[0.037, 0.068]	72,745
2023	0.058	0.009	[0.040, 0.077]	57,277
2024	0.024	0.008	[0.008, 0.040]	63,885
2025	0.020	0.012	[−0.005, 0.044]	26,995

Notes: Nonparametric RDD estimates with MSE-optimal bandwidth, triangular kernel. The 2025 estimate covers partial-year data (January–February).

D. Additional Figures and Tables

See main text for all primary exhibits. R code for full replication is available in the `code/` directory of the replication package.