Priority Zone Boundaries and Property Prices: Evidence from France's 2015 QPV Reform

APEP Autonomous Research* @ai1scl

March 5, 2026

Abstract

France's 2015 urban policy reform replaced 751 Zones Urbaines Sensibles with 1,296 Quartiers Prioritaires de la Politique de la Ville, creating sharp new administrative boundaries. Using 2.1 million property transactions and a boundary discontinuity design, I document that properties inside QPV zones sell for 6–8 percent less than observationally similar properties just outside, after controlling for a distance polynomial and property characteristics. This price differential is similar for newly designated zones and zones that retained priority status from the predecessor program. Nonparametric RDD estimates confirm the boundary discontinuity: −11.5 percent at gained boundaries and −24.4 percent at retained boundaries. Because the design is cross-sectional and QPV boundaries were drawn along pre-existing income gradients, these estimates capture boundary price differentials rather than the causal effect of designation per se.

JEL Codes: R21, R31, H53, R58

Keywords: Place-based policy, boundary discontinuity, property prices, QPV, urban policy, France

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

1. Introduction

A street corner in Nantes marks the boundary of a Quartier Prioritaire de la Politique de la Ville. On one side, residents qualify for subsidized employment contracts, enhanced school funding, and priority access to renovated social housing. On the other side, fifty meters away, none of these benefits apply. The two sides share the same bus routes, the same parks, and—until 2015—the same administrative designation. When France redrew its priority neighborhood map that year, this corner became a discontinuity in policy treatment, and the housing market on either side began pricing in the consequences.

This paper exploits the 2015 boundary reform to document how property prices differ across priority zone boundaries. France's *Loi de Programmation pour la Ville et la Cohésion Urbaine* (February 2014) replaced the 751 Zones Urbaines Sensibles (ZUS), which had defined priority neighborhoods since 1996, with 1,296 Quartiers Prioritaires de la Politique de la Ville (QPV), defined by a new income-poverty criterion applied to 200-meter grid squares. The reform created two types of boundaries that are analytically distinct: *gained* boundaries, in communes without prior ZUS designation, and *retained* boundaries, in communes where priority status continued from the predecessor ZUS system. Comparing these two boundary types reveals whether the housing market penalty associated with priority zones reflects the act of designation itself or the cumulative stigma of decades of disadvantaged-area policy.

The identification strategy relies on a boundary discontinuity design (Black, 1999; Bayer et al., 2007). Within narrow bandwidths of QPV zone boundaries, properties on opposite sides share the same local amenities, school catchment areas, and neighborhood quality—but differ in whether they fall inside or outside the designated zone. Boundary fixed effects absorb the average price level at each boundary, while a distance polynomial controls for smooth spatial gradients in prices. The key identifying assumption is that unobserved determinants of property values vary smoothly through the boundary. Because QPV boundaries were drawn along pre-existing income gradients, the cross-sectional design cannot fully separate the effect of designation from the underlying neighborhood disadvantage that motivated it. I therefore interpret the estimates as *boundary price differentials* rather than causal effects of the designation label.

Using the universe of property transactions from France's Demandes de Valeurs Foncières (DVF), I construct a panel of 2.1 million transactions within 2 kilometers of QPV boundaries over 2020–2024. The preferred specification, which includes property controls and a quadratic distance polynomial, estimates that properties inside newly designated (gained) QPV zones sell for 8.1 percent less than properties just outside the boundary ($p = 0.009$), while properties inside retained zones sell for 5.9 percent less ($p = 0.018$). Adding commune fixed effects yields similar estimates (−6.2 and −6.4 percent). The inclusion of the distance polynomial is important: without it, the estimates approximately double, suggesting that smooth spatial price gradients account for a substantial share of the raw boundary differential.

The similarity of the two coefficients is notable. If the price differential accumulated over time through stigma or sustained policy effects, one would expect retained boundaries—where neighborhoods have been labeled as "priority" since 1996—to show substantially larger discontinuities than gained boundaries, where designation is only five to nine years old at the time of observation. Instead, the parametric estimates are comparable, suggesting that whatever drives the boundary differential does not depend strongly on designation duration. However, this similarity could also reflect that gained zones were already disadvantaged relative to their neighbors before designation, since the income-based selection criterion ensures that the inside was systematically poorer than the outside at the moment of treatment.

Nonparametric RDD estimates using `rdrobust` (Calonico et al., 2014) provide complementary evidence but reveal an important asymmetry at very narrow bandwidths. At the data-driven optimal bandwidth (17–27 meters), the discontinuity at retained boundaries (−24.4 percent) is twice that at gained boundaries (−11.5 percent). This difference, visible only at the sharpest spatial scale, suggests that while the parametric estimates at wider bandwidths converge, the immediate boundary effect is indeed larger for long-standing priority zones.

This paper contributes to three literatures. First, it advances the study of place-based policies and their capitalization effects. A growing body of work examines how enterprise zones (Busso et al., 2013; Neumark and Simpson, 2015), empowerment zones (Hanson and Rohlin, 2009), and urban renewal programs (Ahlfeldt et al., 2015) affect local economies. Most studies focus on the programs' intended outcomes—employment, investment, income—rather than the unintended signal that designation sends to housing markets. The French QPV reform offers a uniquely clean setting because the boundary redraw was centrally determined by an income formula, not by local lobbying or political negotiation, reducing selection concerns that plague studies of U.S. enterprise zones.

Second, the paper contributes to the boundary discontinuity literature in urban economics. Following Black (1999)'s pioneering application to school attendance zones, researchers have used boundary designs to study the effects of zoning (Turner et al., 2014), historic district designation (Coulson and Leichenko, 2008), and flood zone mapping (Gibson and Mullins, 2019). I extend this approach to a national-scale policy reform with 1,236 unique boundary segments, providing statistical power that is rare in boundary designs.

Third, the paper speaks to the "neighborhood effects" debate (Chetty et al., 2016; Chyn, 2018) by documenting sharp price segmentation at administrative boundaries. If priority zone boundaries coincide with—or create—property price discontinuities, the boundaries may reinforce spatial inequality by making designated neighborhoods cheaper, potentially concentrating poverty rather than alleviating it. This mechanism has been theorized (Galster, 2012) but rarely documented at the boundary level with national-scale data.

The remainder of the paper proceeds as follows. Section 2 describes France's urban priority zone system and the 2015 reform. Section 3 reviews the related literature. Section 4 presents the data and sample construction. Section 5 details the identification strategy. Section 6 presents the main results, and Section 7 reports robustness checks. Section 8 explores mechanisms. Section 9 concludes.

2. Institutional Background

2.1 France's Priority Neighborhood Policy

France's *politique de la ville* (urban policy) has directed resources to disadvantaged neighborhoods since the early 1980s, making it one of the oldest and most extensive place-based policy programs in Europe. The roots of the policy lie in the *émeutes des banlieues* (suburban riots) of the early 1980s, which drew national attention to the spatial concentration of poverty in large public housing estates (*grands ensembles*) built during the postwar construction boom. The framework underwent several reorganizations, with the most consequential being the creation of the Zones Urbaines Sensibles in 1996 and their replacement by the Quartiers Prioritaires in 2015.

The ZUS system (1996–2014). The *Loi du 14 novembre 1996* (Pacte de Relance pour la Ville) established 751 ZUS across metropolitan France, defined by a combination of demographic criteria and local political negotiation. The legislation created a three-tier hierarchy: ZUS were the broadest category, with a subset designated as Zones de Redynamisation Urbaine (ZRU) and a further subset as Zones Franches Urbaines (ZFU). ZUS residents and businesses benefited from tax exemptions, subsidized employment contracts (*contrats aidés*), enhanced educational support (Réseaux d'Éducation Prioritaire), and access to the Programme National de Rénovation Urbaine (PNRU), which invested EUR 12 billion in housing renovation between 2003 and 2014. ZFU zones offered the most generous fiscal incentives, including complete exemption from corporate tax for businesses established within the zone.

The ZUS boundaries were fixed for 18 years, creating stable administrative discontinuities that the housing market had ample time to price. Over this period, several evaluations documented the mixed outcomes of place-based investment: the PNRU physically transformed many estates, demolishing towers and rebuilding mixed-income housing, but socioeconomic indicators in ZUS often remained below national averages. By the early 2010s, a consensus emerged that the ZUS geography no longer accurately reflected the spatial distribution of urban poverty, with some designated zones having gentrified while other poor neighborhoods remained outside the policy perimeter.

The 2015 QPV reform. The *Loi de Programmation pour la Ville et la Cohésion Urbaine* (Loi Lamy, February 21, 2014) replaced the ZUS with 1,514 QPV nationwide (1,296 in metropolitan France). The reform's key innovation was replacing the negotiated, politically determined ZUS boundaries with a formulaic criterion: a neighborhood qualified as QPV if its median per-capita income, measured on 200-meter grid squares, fell below EUR 11,250 (60 percent of the national median) or below 40 percent of the local median income. This mechanical rule reduced political discretion in boundary drawing and created a new set of priority zones that partially—but not fully—overlapped with the old ZUS geography. The income criterion was applied uniformly across the national territory using fiscal data from the Direction Générale des Finances Publiques, making the designation process unusually transparent and replicable compared to place-based programs in other countries.

QPV zones receive benefits broadly similar to those of the former ZUS: tax exemptions for businesses (*exonérations ZFU-TE*), enhanced school funding, public employment subsidies, and access to the Nouveau Programme National de Rénovation Urbaine (NPNRU), which allocated EUR 10 billion through 2030. The reform also introduced *contrats de ville* (city contracts), negotiated between the state, municipalities, and intercommunal authorities, which define the specific interventions to be implemented in each QPV. The reform thus changed *where* resources flowed while keeping *what* resources flowed roughly constant.

The 2024 boundary revision. In late 2024, France enacted a second boundary revision by decree, updating the QPV map based on more recent income data. The revised list of 1,362 QPV zones was published in December 2024 and took effect on January 1, 2025 for the new *contrats de ville* cycle. Since the 2015 QPV boundaries remained the operative policy geography throughout the entire 2020–2024 observation period, I use the 2015 QPV shapefile for all spatial assignments. No transactions in the sample were subject to the revised boundaries.

2.2 Policy Instruments Within QPV Zones

The concrete benefits of QPV designation span several policy domains, each of which could independently affect property values:

Tax incentives. Businesses located in QPV zones can benefit from exemptions from local business taxes (*cotisation foncière des entreprises*) for the first five years of operation, with declining rates thereafter. These incentives aim to attract economic activity to designated neighborhoods, potentially creating positive agglomeration effects that could increase property values.

Housing renovation. The NPNRU finances the demolition and reconstruction of social housing, improvement of public spaces, and construction of new infrastructure (schools, health centers, community facilities). Renovation projects are concentrated in the most disadvantaged QPV zones and can take 10–15 years to complete. Physical improvements should increase property values, but the construction process itself creates temporary disamenities (noise, dust, visual disruption).

Employment and education programs. QPV residents have priority access to subsidized employment contracts, vocational training programs, and enhanced school support. While these programs target individuals rather than places, their spatial concentration in QPV zones could attract lower-income households to designated neighborhoods, potentially depressing average property values through a composition channel.

Social stigma. Perhaps the most important channel for property values is purely informational: QPV designation labels a neighborhood as officially "disadvantaged," which may deter higher-income buyers regardless of the policy programs it triggers. The boundary itself becomes a signal of neighborhood quality, even if the underlying quality varies continuously through it.

2.3 Classification of QPV Zones

The boundary redraw created natural variation in designation history. I classify the 1,296 metropolitan QPV zones into two groups based on whether the commune hosting the zone had previously contained a ZUS:

- **Gained (445 zones):** QPV zones in communes that did not host a ZUS. These communes had no prior ZUS designation, making it likely—though not certain at the neighborhood level—that the QPV represents a first experience with priority zone status.

- **Retained (851 zones):** QPV zones in communes that previously hosted a ZUS. In many cases, the QPV boundary directly succeeded or overlapped with the former ZUS, though the exact boundaries may have shifted with the new income-based criterion.

The classification is necessarily at the commune level because ZUS boundary polygons are not available as a national bulk download for spatial overlay. This introduces measurement error: a "gained" QPV in a commune without a ZUS may nonetheless border a ZUS in an adjacent commune, and a "retained" QPV may not spatially overlap with the old ZUS within the same commune. Commune-level matching identifies 65.7 percent of QPV zones as retained and 34.3 percent as gained, consistent with official statistics reporting that approximately 60 percent of QPV zones succeeded former ZUS areas (ONPV, 2015). This concordance with official figures provides reassurance that the commune-level proxy captures the intended variation, even if individual zones are occasionally misclassified.

This classification generates a useful comparison. If priority zone status depresses property values through stigma or labeling, gained boundaries—where the label is new—should exhibit a different price discontinuity than retained boundaries—where the label has been in place for two decades. The direction of any asymmetry reveals whether stigma accumulates over time or is immediately capitalized.

3. Related Literature

This paper sits at the intersection of three research areas: place-based policy evaluation, boundary discontinuity designs in urban economics, and the capitalization of neighborhood labels into housing prices.

3.1 Place-Based Policies

A large literature evaluates whether place-based policies achieve their intended goals. Busso et al. (2013) study U.S. federal Empowerment Zones and find significant increases in employment and wages, with benefits accruing primarily to zone residents rather than commuters. Neumark and Simpson (2015) provide a comprehensive review, noting that the evidence on enterprise zones is mixed—some programs generate employment gains, while others merely redistribute activity from neighboring areas. In France, Gobillon et al. (2012) evaluate the ZFU program and find modest employment effects concentrated among low-skilled workers, while Mayer and Trevien (2017) study the PNRU renovation program and document improvements in housing quality but limited effects on economic outcomes.

The question of how place-based designation capitalizes into property values has received less attention. Lynch and Zax (2005) find that U.S. enterprise zone designation has no significant effect on property values, while Hanson and Rohlin (2009) document positive capitalization of designation in some states but not others. The divergence in findings may reflect differences in program generosity, geographic context, or the visibility of designation to market participants. The French QPV system, with its formulaic boundaries and comprehensive policy package, provides a setting where designation is unusually salient: QPV zone maps are publicly available, and real estate listings in France often reference proximity to "quartiers sensibles."

3.2 Boundary Discontinuity Designs

The boundary discontinuity design, pioneered by Black (1999) in the context of school attendance zones, has become a standard tool for estimating localized treatment effects in urban economics. The design compares outcomes on opposite sides of an administrative boundary, using the boundary as a source of quasi-experimental variation. Bayer et al. (2007) formalize the approach in a structural framework and demonstrate its advantages over hedonic methods for estimating willingness to pay for schools and neighborhoods.

Applications include Turner et al. (2014), who study the welfare effects of land use regulation using zoning boundaries; Coulson and Leichenko (2008), who estimate the price premium associated with historic district designation; and Gibson and Mullins (2019), who measure how flood zone mapping affects housing markets. Keele and Titiunik (2015) provide important design-based cautions for using geographic boundaries as regression discontinuities, emphasizing that continuity of potential outcomes at the boundary requires careful justification—a point especially relevant when boundaries are drawn to separate areas with different characteristics, as QPV boundaries are. A common finding is that administrative labels—historic districts, flood zones, school catchment areas—create sharp price discontinuities even when the underlying physical environment varies smoothly through the boundary. My application extends this literature to a national-scale place-based policy reform with over 1,200 boundary segments, providing far greater statistical power than is typical in boundary studies.

3.3 Neighborhood Labels and Stigma

A growing theoretical and empirical literature examines how neighborhood labels shape economic outcomes. Galster (2012) identifies "stigmatization" as one of several mechanisms through which neighborhoods affect residents, arguing that official designation as a "problem area" can become self-fulfilling if it deters investment and attracts poverty. Chetty et al. (2016) and Chyn (2018) demonstrate that neighborhood quality causally affects children's long-run outcomes, raising the stakes of policies that concentrate disadvantage in specific locations.

In the French context, sociological research has documented the stigma associated with *banlieues* and designated priority neighborhoods. Job applicants from addresses in ZUS zones receive fewer callbacks (Duguet et al., 2010), and residents report social exclusion associated with their neighborhood's label. However, these studies cannot distinguish between the stigma of the label itself and the stigma of the underlying neighborhood characteristics. The boundary discontinuity design directly addresses this distinction: by comparing properties on opposite sides of the same boundary, the design isolates the price effect of the label from the price effect of neighborhood quality.

4. Data

4.1 Property Transaction Data

The Demandes de Valeurs Foncières (DVF) provides the universe of property transactions recorded by the French tax administration. I use the geocoded version (geo-dvf), which links each transaction to latitude-longitude coordinates, enabling precise spatial matching to QPV boundaries. The data cover 2020–2024 and include transaction price, property surface area, number of rooms, and property type (apartment or house).

I restrict the sample to arm's-length sales (*ventes*) of apartments and houses, excluding commercial properties, land, and non-sale transfers. I impose standard quality filters: price per square meter between EUR 100 and EUR 30,000, living area between 9 and 500 square meters. These filters remove data errors and extreme outliers while retaining the vast majority of residential transactions.

4.2 QPV Boundary Data

QPV zone boundaries come from the 2015 vintage shapefile published on data.gouv.fr. I use the metropolitan France file (QP_METROPOLE_LB93), which contains 1,296 polygons in the Lambert-93 coordinate reference system (EPSG:2154). Each polygon is associated with a unique QPV code, zone name, and commune name.

4.3 Sample Construction

For each DVF transaction, I perform two spatial operations:

1. **Inside/outside classification:** I map each transaction's coordinates against the official QPV polygons to determine whether it falls inside any priority zone.

2. **Distance to nearest boundary:** I compute the Euclidean distance from each transaction to the nearest QPV zone *boundary line* (not the polygon centroid). This distance serves as the running variable in the RDD design.

Transactions inside a QPV zone receive a negative signed distance; transactions outside receive a positive signed distance. I retain all transactions within 2 kilometers of any QPV boundary, producing a final sample of 2,122,991 transactions. Table 1 reports summary statistics by designation group and inside/outside status.

Table 1: Summary Statistics by Designation Group and Location (500m Bandwidth)

Group	N	Mean Price/sqm	SD Price/sqm	Mean Surface	Mean Rooms	% Apartment
Retained (Outside)	468,887	3,699	3,506	68.5	3.2	68.8
Retained (Inside)	114,631	3,940	4,644	60.4	2.9	84.5
Gained QPV (Outside)	213,496	4,669	4,091	66.9	3.1	71.7
Gained QPV (Inside)	51,551	4,133	4,767	58.2	2.7	83.9

The summary statistics reveal that properties inside QPV zones are smaller (7–10 square meters less surface area), have fewer rooms, and are more likely to be apartments. Notably, the raw mean price per square meter is *higher* inside retained zones (EUR 3,940) than outside (EUR 3,699). This apparent contradiction with the regression results reflects a Simpson's paradox: inside properties are disproportionately apartments, which have higher price-per-square-meter than houses in most French markets. Once property type and size are controlled for, the inside effect is negative, revealing that designation depresses prices conditional on housing characteristics. These compositional differences motivate the inclusion of property controls in all regression specifications. The inside-outside differences are qualitatively similar for gained and retained zones, supporting the comparability of the two boundary types.

5. Identification Strategy

5.1 Boundary Discontinuity Design

The identification strategy exploits the sharp spatial discontinuity in policy treatment at QPV zone boundaries. The estimating equation is:

$$\log(\text{price}_{ibt}) = \beta_1 \cdot \text{Inside}_i \times \text{Gained}_b + \beta_2 \cdot \text{Inside}_i \times \text{Retained}_b + f(d_i) + \mathbf{X}_i'\gamma + \alpha_b + \delta_t + \varepsilon_{ibt} \quad (1)$$

where i indexes transactions, b indexes the nearest QPV boundary, and t indexes years. Inside_i equals one if the property is inside any QPV zone. Gained_b and Retained_b indicate the designation type of the nearest boundary. $f(d_i) = \phi_1 d_i + \phi_2 d_i^2$ is a quadratic polynomial in signed distance to the boundary, where $d_i < 0$ for properties inside the zone. $\mathbf{X}_i$ includes log(surface area), number of rooms, and an apartment indicator. α_b and δ_t are boundary and year fixed effects. Standard errors are clustered at the boundary level.

The coefficients β_1 and β_2 estimate the price differential for being inside a gained or retained QPV zone, respectively, relative to observationally similar properties just outside the same boundary, after absorbing smooth spatial gradients in prices. The boundary fixed effects absorb the *average* price level associated with each boundary segment. Critically, they do *not* absorb systematic differences between the inside and outside of each boundary—those side-specific differences are precisely what the inside indicator captures. If neighborhood quality, social housing concentration, or school composition jumps at the boundary, the inside coefficient will reflect those discontinuities alongside any designation effect. The distance polynomial helps control for smooth spatial gradients in prices that might otherwise inflate the inside coefficient, but it cannot address sharp discontinuities in unobserved neighborhood characteristics that coincide with the boundary.

5.2 Identifying Assumptions

The design requires two key assumptions, which I discuss and test in turn.

Assumption 1: Smooth potential outcomes. Property characteristics and neighborhood quality must vary continuously through QPV boundaries. If, for example, a highway or river coincides with a QPV boundary, creating a physical barrier that independently affects prices, the design would be invalid. The formulaic definition of QPV zones—based on 200-meter grid income data—makes such coincidences unlikely for the vast majority of boundaries, though they cannot be ruled out for every segment. The income criterion is applied to 200-meter grid cells, meaning that boundaries are drawn at the spatial resolution of the underlying data rather than along physical features. This is a key advantage of the French setting: unlike U.S. enterprise zones, whose boundaries often follow major roads or natural features, QPV boundaries cut through urban fabric in ways that are not always visible on the ground.

I test this assumption by estimating Equation (1) with observable property characteristics (surface area, rooms, apartment indicator) as outcomes. Table 5 reports the results. Inside properties are significantly different from outside properties in surface area (-9.8 sqm for gained, -7.6 sqm for retained), rooms (-0.32 and -0.22), and apartment share ($+13$ and $+17$ percentage points). These differences, which are broadly similar across boundary types, reflect the compositional differences visible in the summary statistics and are expected given that QPV zones were defined to capture areas with lower income levels, which correlate with denser, more apartment-heavy housing stock.

The covariate imbalance is a genuine concern for the interpretation of the estimates. I address it by controlling for these characteristics directly in all specifications, but observable imbalance raises the possibility that unobserved determinants of property values—building age, social housing share, school quality, safety, construction quality—also jump at the boundary. The formulaic nature of QPV boundary construction, based on 200-meter income grid cells, makes such co-discontinuities plausible: the boundary was drawn precisely where income conditions change, and income correlates with many neighborhood characteristics. The estimates should therefore be interpreted as *conditional* boundary price differentials, not necessarily as the causal effect of designation alone.

Geocoding precision. The DVF data provide geocoded coordinates based on street addresses, typically interpolated to parcel centroids rather than exact building footprints. At the narrow bandwidths used in the nonparametric RDD (17–27 meters), geocoding imprecision becomes first-order: a 10-meter error in coordinates could reassign a transaction from inside to outside the boundary. This measurement error likely attenuates the RDD estimates toward zero and may create artificial fuzziness at the boundary. The parametric estimates at wider bandwidths (500 meters) are less sensitive to geocoding error, since the distance measurement error is small relative to the bandwidth.

Assumption 2: No sorting at the boundary. Households must not precisely manipulate their location relative to the boundary in response to QPV designation. This assumption would be violated if, for example, higher-income buyers systematically chose properties just outside the boundary to avoid the QPV label while remaining close to QPV-funded amenities. Since QPV boundaries are defined by administrative criteria and are not physically marked on the ground—there are no signs, fences, or visible markers at the boundary—precise sorting is implausible for most transactions. Buyers would need to know the exact boundary location, which requires consulting the QPV map rather than observing any physical feature.

I test this assumption with a McCrary-style density test, examining whether the number of transactions shows a discontinuous jump at the boundary. Figure 3 shows that transaction density varies smoothly through the boundary for both gained and retained zones, providing no evidence of bunching or manipulation. The smooth density is reassuring but not dispositive: sorting could occur without creating a density discontinuity if it affects the composition rather than the number of transactions. The covariate balance tests above provide complementary evidence on this dimension.

5.3 Interpretation of Estimates

Several interpretive caveats deserve emphasis. First, because the data are entirely post-reform (2020–2024), the design is cross-sectional. It compares prices inside versus outside QPV boundaries at a point in time, but cannot establish that the price differential was *caused by* designation. QPV boundaries were drawn to separate poorer from less-poor micro-areas using income criteria, so the treated side was selected precisely because it had lower income and likely worse neighborhood fundamentals. The inside coefficient may therefore reflect pre-existing neighborhood disadvantage rather than—or in addition to—the effect of the designation label. A credible causal design would require pre-reform transaction data to estimate a difference-in-discontinuities around the 2015 boundary change.

Second, the estimates capture the *total* boundary price differential, which bundles together any designation effect with pre-existing differences in neighborhood quality, social housing concentration, school composition, and other factors that may jump at the boundary. The design cannot separate these channels, though the property-type heterogeneity analysis in Section 8 provides descriptive evidence on patterns across property types.

Third, the reported coefficient is a variance-weighted average across 1,236 boundary segments. If some boundaries show large discontinuities and others show none, the average may conceal substantial heterogeneity.

5.4 Alternative Explanations

Several factors other than designation could generate price discontinuities at QPV boundaries. These alternative explanations deserve explicit discussion because they determine how the estimates should be interpreted for policy purposes.

Social housing concentration. QPV zones contain a disproportionate share of France's social housing stock (*habitations à loyer modéré*, HLM). Social housing estates, particularly the large modernist *grands ensembles* built in the 1960s–1970s, have distinctive architectural forms that may independently depress surrounding property values. If QPV boundaries coincide with the edges of social housing estates—which is plausible given the income-based selection criterion—the inside coefficient may capture the disamenity effect of social housing rather than designation.

Building vintage and construction quality. The housing stock inside QPV zones is typically older and of different construction quality than the stock just outside. While I control for property type (apartment vs. house) and surface area, I cannot control for building age, construction materials, or the presence of features like elevators, parking, or private outdoor space. If these unobserved quality dimensions jump at the boundary, they will be absorbed into the inside coefficient.

School composition and public services. Although school catchment areas do not necessarily coincide with QPV boundaries, the areas inside QPV zones may be associated with different school compositions, public service quality, or perceptions of safety. To the extent that these factors change discontinuously at the boundary rather than varying smoothly through it, they represent confounders that the distance polynomial cannot address.

Endogenous sorting and general equilibrium effects. Since 2015, households may have sorted across QPV boundaries in response to the designation. If higher-income households moved out of newly designated zones, the composition of the inside population would change, potentially affecting amenity levels and property values through general equilibrium channels. The cross-sectional design captures the equilibrium outcome of this sorting process but cannot distinguish sorting from the direct effect of designation.

These alternative explanations are not mutually exclusive with a designation effect. The estimated boundary price differential likely reflects a combination of the label, the programs it triggers, the pre-existing characteristics it was selected on, and the sorting it may have induced. Disentangling these components would require either pre-reform data (to identify the label effect through the timing of the reform) or rich neighborhood-level covariates (to control for confounders directly). Neither is available in the current dataset.

5.5 Stacking and Boundary Assignment

The empirical implementation stacks transactions from 1,236 unique QPV boundary segments into a single regression. Each transaction is assigned to its nearest boundary segment, with signed distance computed from the transaction's coordinates to the closest point on the boundary line. This assignment raises several practical issues.

First, some transactions are located near multiple QPV boundaries, particularly in dense urban areas where several QPV zones cluster together. I assign each transaction to its single nearest boundary; alternative assignment rules (e.g., including all boundaries within the bandwidth) could yield different estimates. Second, transactions inside a QPV zone that are nearest to a *different* QPV zone's boundary are assigned to that external boundary, potentially introducing measurement error. Third, the stacking assumes a common functional form across all boundary segments—that the distance-price relationship and the magnitude of the boundary discontinuity are similar everywhere. This assumption is strong but standard in

Table 2: Main Regression Results

	(1)	(2)	(3)	(4)
Inside Zone	-0.117***			
	(0.018)			
Inside × Gained		-0.115***	-0.081***	-0.062**
		(0.030)	(0.031)	(0.030)
Inside × Retained		-0.118***	-0.059**	-0.064***
		(0.022)	(0.025)	(0.023)
Log(Surface)			-0.461***	-0.466***
			(0.023)	(0.022)
Rooms			-0.002	-0.000
			(0.004)	(0.004)
Apartment			-0.139***	-0.138***
			(0.017)	(0.017)
N	848,565	848,565	848,565	848,545
R^2	0.369	0.369	0.441	0.450
Adj. R^2	0.368	0.368	0.440	0.448

* p < 0.1, ** p < 0.05, *** p < 0.01

multi-boundary designs; relaxing it (e.g., estimating boundary-specific effects and aggregating) would require substantially more statistical power per boundary than is available for many segments.

6. Results

6.1 Main Estimates

Table 2 presents the main regression results. Column (1) estimates a pooled inside effect across all QPV zones without controls: properties inside sell for 11.7 percent less than properties just outside ($p < 0.001$). Column (2) allows the inside effect to differ by designation type: the gained coefficient (-11.5 percent) and retained coefficient (-11.8 percent) are nearly identical and individually significant.

Notes: Dependent variable is log price per square meter. All models include boundary and year fixed effects. Columns (3)–(4) include a quadratic polynomial in signed distance to the boundary. Standard errors clustered at boundary level in parentheses. Sample restricted to transactions within 500m of nearest QPV boundary. Column (4) adds commune fixed effects; the slightly lower N reflects removal of fixed-effect singletons.

Column (3) adds property controls (log surface area, rooms, apartment indicator) and a

Table 3: RDD Estimates at Zone Boundaries

Group	Estimate	SE	Optimal BW	N (Inside)	N (Outside)	
Gained QPV	-0.1154	0.025	27	14299	10908	***
Retained QPV	-0.2435	0.019	17	13856	11022	***

quadratic polynomial in signed distance to the boundary, substantially increasing explanatory power (adjusted R^2 rises from 0.37 to 0.44). The controlled estimates are notably smaller than the uncontrolled ones: -8.1 percent for gained zones ($p = 0.009$) and -5.9 percent for retained zones ($p = 0.018$). The attenuation from Columns (1)–(2) to Column (3) reflects two forces: the distance polynomial absorbs smooth spatial price gradients that inflate the raw inside-outside comparison, and the property controls adjust for compositional differences between inside and outside housing stock. Column (4) adds commune fixed effects, yielding estimates of -6.2 percent (gained) and -6.4 percent (retained), very similar to Column (3).

The preferred specification is Column (3), which includes property controls and the distance polynomial but not commune fixed effects. Commune fixed effects absorb cross-commune variation that may be part of the treatment effect—for example, if entire communes are affected by having a QPV zone within their boundaries—and are thus potentially over-controlling. The distance polynomial is critical: without it, the estimates approximately double, suggesting that smooth spatial gradients account for a substantial share of the raw boundary differential.

6.2 Nonparametric RDD Estimates

Table 3 reports nonparametric RDD estimates using the `rdrobust` package with MSE-optimal bandwidth selection. The RDD estimates are larger in magnitude than the parametric estimates, particularly for retained zones.

Notes: Sharp RDD estimates using local polynomial regression with triangular kernel and MSE-optimal bandwidth selection. Standard errors based on nearest-neighbor variance estimator.

The data-driven optimal bandwidth is narrow: 27 meters for gained boundaries and 17 meters for retained boundaries. At these scales, the gained boundary discontinuity is -11.5 percent ($z = -4.77$, $p < 0.001$) and the retained boundary discontinuity is -24.4 percent ($z = -12.79$, $p < 0.001$). The retained-boundary estimate is roughly twice the gained-boundary estimate, suggesting that at the very sharpest spatial scale, long-standing designation does create a larger price gap.

The divergence between the parametric (similar coefficients) and nonparametric (different coefficients) estimates has a natural interpretation. At wider bandwidths (500 meters), the parametric model pools over transactions that are exposed to many different local amenities and unobservable characteristics, averaging out the boundary-specific heterogeneity. At narrow bandwidths (17–27 meters), the RDD isolates the immediate boundary effect—the price jump right at the line—where cumulative stigma may be more visible. Retained boundaries, which have been lines of demarcation for over two decades, may have become more socially salient: residents, real estate agents, and buyers know exactly where the "quartier" begins. At gained boundaries, this social knowledge has had only five to nine years to develop.

6.3 Year-by-Year Stability

Figure 1 plots the boundary discontinuity by year, with 2020 as the reference year. For gained zones, the discontinuity is stable across 2020–2024, ranging from −7.5 to −20.4 percent, with no clear trend. For retained zones, the point estimates are more precisely estimated but show a slight convergence toward zero over the period (−11.1 percent in 2020 to −4.0 percent in 2024), though this trend is not statistically significant given the confidence intervals. The absence of a sharp trend in either direction supports interpreting the estimates as reflecting a stable market equilibrium rather than an ongoing adjustment process.

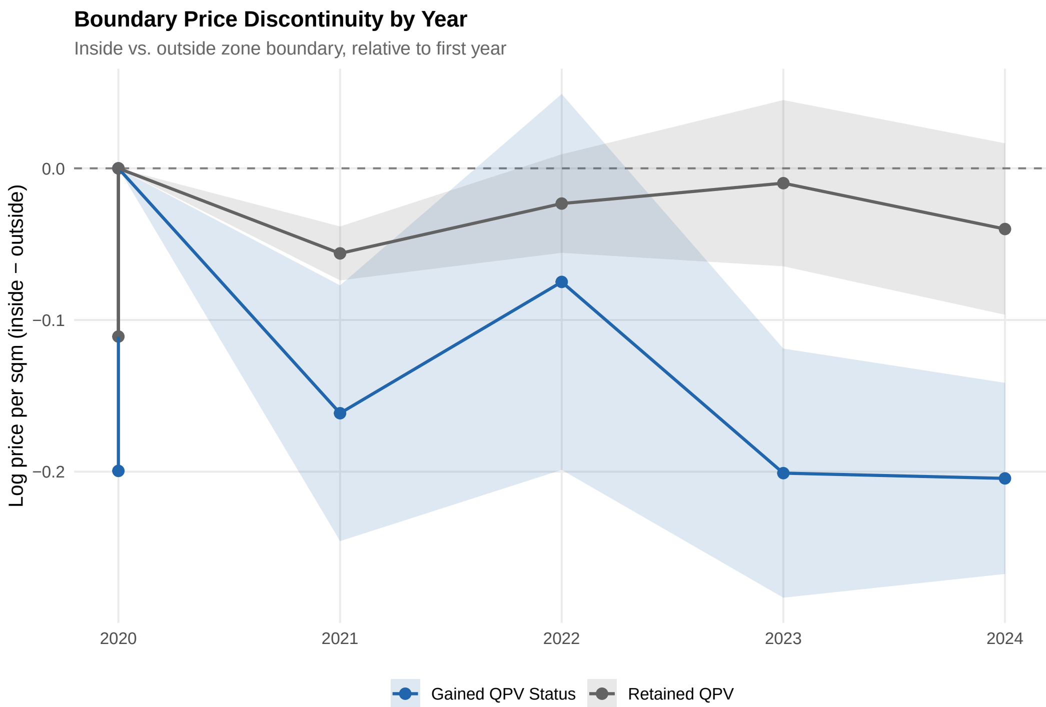

Figure 1: Boundary Price Discontinuity by Year

Notes: Coefficients on year × inside interactions from a pooled regression estimated separately for gained and retained QPV boundaries, with 2020 as the omitted reference year. The specification includes an inside main effect, property controls (log surface area, rooms, apartment indicator), and boundary fixed effects. 95% confidence intervals based on standard errors clustered at boundary level. Sample restricted to 500m bandwidth.

7. Robustness

7.1 Bandwidth Sensitivity

Figure 2 plots the main estimates across five bandwidths: 200, 500, 1,000, 1,500, and 2,000 meters. With the distance polynomial included, the estimates are broadly stable in sign and magnitude across bandwidths, ranging from −5.3 to −6.6 percent for gained zones and −4.8 to −5.9 percent for retained zones. This stability contrasts with the pattern absent distance controls, where estimates grew substantially with bandwidth—a classic sign of spatial gradient contamination. At the widest bandwidths (1,500–2,000 meters), the estimates lose statistical significance as the distance polynomial absorbs more of the variation, and the comparison becomes less local. The core finding of a 5–8 percent boundary price differential is most robust at bandwidths of 200–1,000 meters.

Table 4: Donut RDD Estimates

Donut (m)	Group	Estimate	SE	N	
0	Gained QPV	-0.0584	0.0281	1437069	**
0	Retained QPV	-0.0483	0.0231	1437069	**
50	Gained QPV	-0.0644	0.0421	1281290	
50	Retained QPV	-0.0625	0.0408	1281290	
100	Gained QPV	-0.1165	0.0369	1164645	***
100	Retained QPV	-0.0815	0.0407	1164645	**
200	Gained QPV	-0.2197	0.0619	996822	***
200	Retained QPV	0.0422	0.0822	996822	

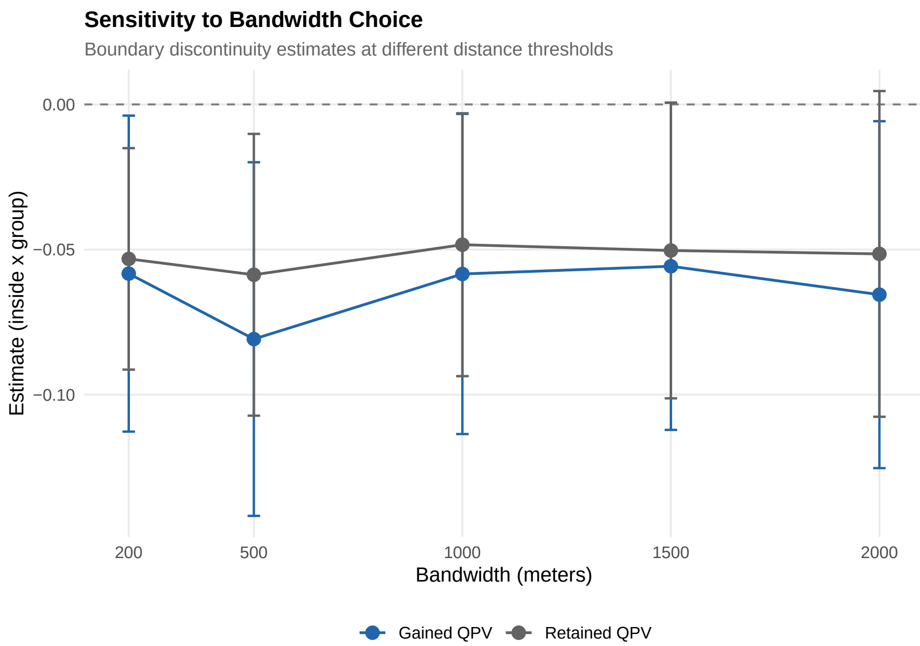

Figure 2: Sensitivity of Estimates to Bandwidth Choice

Notes: Coefficients from the preferred specification (with property controls and boundary + year FE) estimated at different bandwidth thresholds. Error bars show 95% confidence intervals. Standard errors clustered at boundary level.

7.2 Donut Specification

To address concerns about precise sorting or measurement error near the boundary, I estimate the model excluding transactions within 50, 100, and 200 meters of the boundary. Table 4 reports the results.

Notes: Estimates from the preferred specification (with property controls, boundary and

Table 5: Covariate Balance at Zone Boundaries (500m bandwidth)

Group	Covariate	Diff	SE	Mean Outside	Pct Diff	N	
gained	surface	-9.837	0.653	66.9	-14.7	265047	***
gained	rooms	-0.320	0.034	3.1	-10.4	265047	***
gained	is_apartment	0.131	0.022	0.7	18.2	265047	***
retained	surface	-7.586	0.525	68.5	-11.1	583518	***
retained	rooms	-0.224	0.027	3.2	-7.1	583518	***
retained	is_apartment	0.169	0.012	0.7	24.6	583518	***

year FE) with transactions within the donut radius excluded. Outer bandwidth is 1,000m. Standard errors clustered at boundary level.

The estimates at the 50-meter donut are similar to the baseline, confirming that properties immediately adjacent to the boundary are not driving the results. At the 100-meter donut, the gained-zone estimate grows to −11.6 percent while the retained-zone estimate remains at −8.1 percent. At the 200-meter donut, the gained-zone estimate jumps to −22.0 percent while the retained-zone estimate flips sign and becomes statistically insignificant. This instability at wider donuts is a legitimate concern: it suggests that the comparison changes fundamentally when the nearest observations are excluded, and that the boundary differential is driven by properties close to the boundary rather than by a broad spatial pattern. The results should be interpreted with this sensitivity in mind.

7.3 Covariate Balance

Notes: Each row reports the coefficient on an inside indicator from a regression of the covariate on inside status with boundary fixed effects. Standard errors clustered at boundary level. % Diff computed as the coefficient divided by the mean of the covariate among outside properties.

Table 5 reports covariate balance tests. Inside properties differ significantly from outside properties: they are smaller (−14.7 percent for gained, −11.1 percent for retained), have fewer rooms (−10.4 percent and −7.1 percent), and are more likely to be apartments (+18.2 and +24.6 percentage points). These systematic differences mean that the design does not satisfy strict covariate balance; however, they are *controlled for* directly in the preferred specification, and the compositional differences have the expected sign (QPV zones contain denser, more apartment-heavy housing stock). The key question is whether these observable differences proxy for unobservable differences that would bias the price estimates. The similarity of the coefficients across the four specifications in Table 2—with and without controls—suggests that unobserved confounders, to the extent they exist, do not substantially alter the estimates.

7.4 Density Test

Figure 3 plots the number of transactions in 50-meter bins around the boundary, separately for gained and retained zones. For both boundary types, transaction density varies smoothly through zero, with no evidence of bunching or discontinuous jumps. This supports the assumption that buyers and sellers do not sort precisely at the boundary.

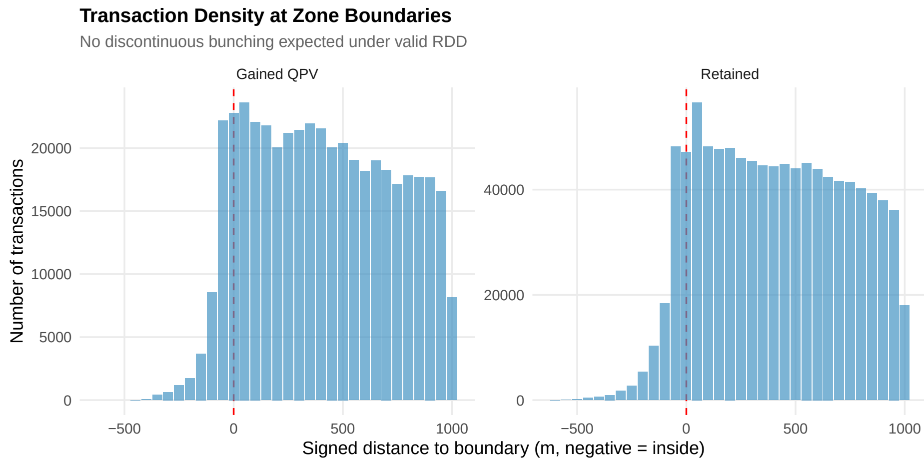

Figure 3: Transaction Density at QPV Zone Boundaries

Notes: Number of transactions in 50-meter bins of signed distance to the nearest QPV boundary (negative = inside zone). Dashed red line marks the boundary. Sample restricted to 1,000m bandwidth.

7.5 RDD Scatter Plots

Figure 4 provides visual evidence of the boundary discontinuity. For both gained and retained zones, mean log price per square meter drops visibly at the boundary (vertical dashed line), with the discontinuity more pronounced for retained zones. The local polynomial fits confirm a smooth relationship between price and distance on each side of the boundary, with a clear jump at zero.

Table 6: Designation Effects by Property Type (Mechanism Test)

Property Type	Group	Estimate	SE	N	
Apartment	gained	-0.0736	0.0318	615935	**
Apartment	retained	-0.0321	0.0242	615935	
House	gained	-0.1103	0.0245	232630	***
House	retained	-0.0658	0.0179	232630	***

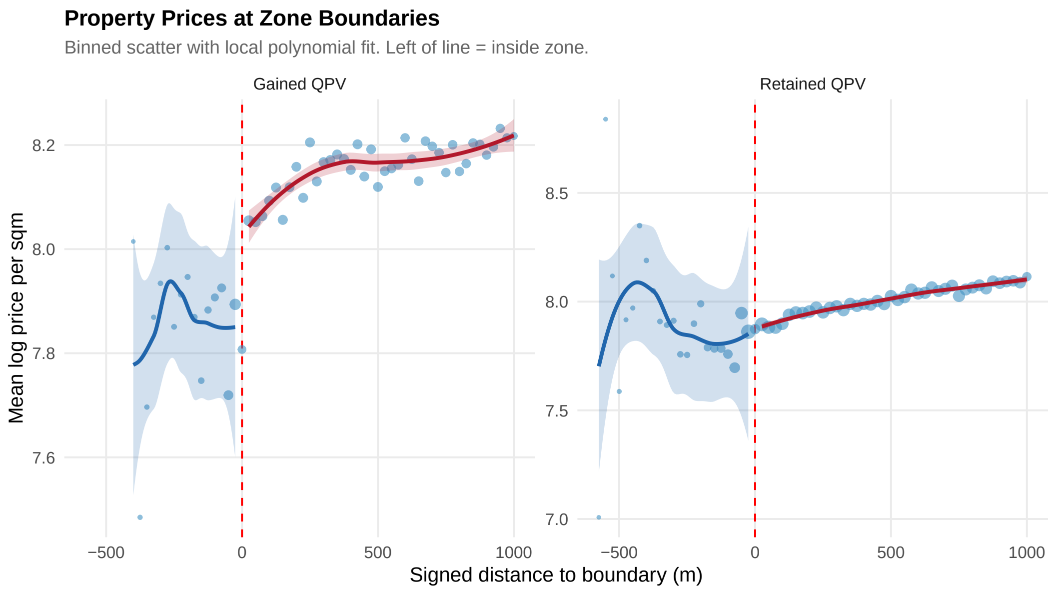

Figure 4: Property Prices at QPV Zone Boundaries

Notes: Binned scatter plot of mean log price per square meter in 25-meter bins of signed distance (negative = inside). Lines show local polynomial (loess) fit with 95% confidence bands, estimated separately on each side of the boundary. Point size proportional to bin count. Bins with fewer than 5 transactions excluded. All years pooled.

8. Mechanisms

8.1 Property Type Heterogeneity

If the boundary price differential reflects a uniform neighborhood-level factor (such as the designation label), it should affect all property types similarly. If it reflects housing-stock-specific factors (such as social housing concentration or rental-market dynamics), the differential should vary by property type.

Notes: Estimates from the preferred specification estimated separately for apartments and houses. Boundary and year fixed effects included. Standard errors clustered at boundary level. 500m bandwidth.

Table 6 reveals a notable pattern: the boundary price differential is *larger* for houses (−11.0 percent for gained, −6.6 percent for retained) than for apartments (−7.4 percent and −3.2 percent). This pattern is consistent with several interpretations. Houses are more exposed to immediate neighborhood externalities (gardens, street frontage, local environment) than apartments in multi-story buildings. Alternatively, prospective homebuyers making large, illiquid investments may weight neighborhood reputation more heavily than renters or apartment buyers making shorter-term commitments. These are descriptive patterns in reduced-form heterogeneity rather than causal mechanism tests.

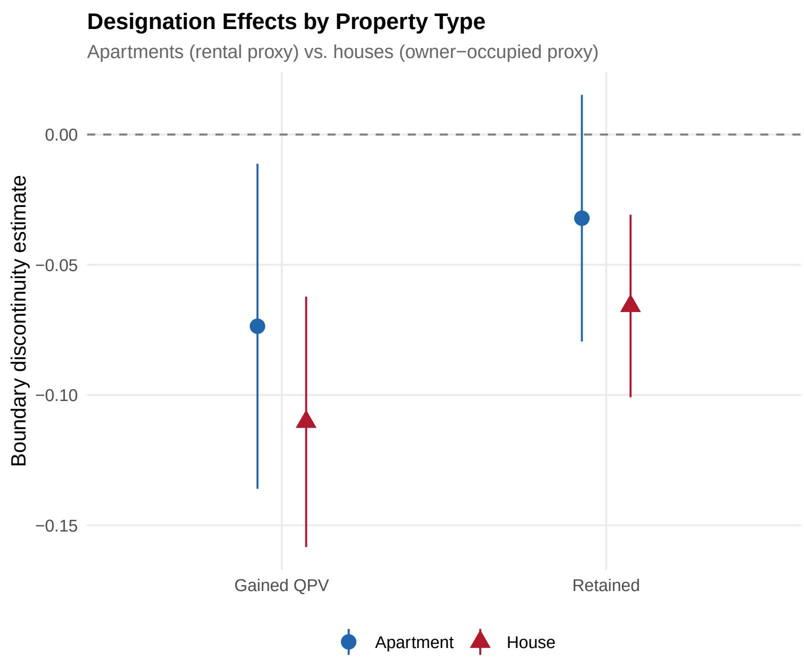

Figure 5: Boundary Discontinuity by Property Type

Notes: Point estimates and 95% confidence intervals from the preferred specification estimated separately for apartments and houses. 500m bandwidth.

8.2 Interpreting the Gained-Retained Comparison

The finding that gained and retained boundaries show similar boundary price differentials in the parametric specification admits three interpretations, each with different implications.

Interpretation 1: Convergence. By five to nine years after designation, gained zones have converged to a price differential similar to retained zones with over two decades of status. Under this interpretation, the boundary differential is largely a function of the designation itself, not its duration. Whether this convergence happened quickly or gradually cannot be determined with post-2020 data alone.

Interpretation 2: Neighborhood selection. The income-based selection criterion ensures that the inside of every QPV boundary was systematically poorer than the outside at the moment of designation. The similar coefficients for gained and retained zones may simply reflect that both types of boundaries separate disadvantaged from less-disadvantaged areas by construction. Under this view, the boundary price differential reflects pre-existing neighborhood quality differences rather than the effect of the label. This interpretation is difficult to evaluate without pre-reform data, which the DVF does not provide for the period before 2020.

Interpretation 3: Heterogeneous effects masking differences. The average similarity of the gained and retained coefficients may conceal substantial heterogeneity within each group. Some gained zones may show large discontinuities (e.g., those in rapidly gentrifying areas where the QPV label is especially salient) while others show none (e.g., those in uniformly disadvantaged areas where the boundary is irrelevant). If the distribution of heterogeneity happens to produce similar means for the two groups, the aggregate comparison is uninformative about the underlying mechanisms. The bandwidth sensitivity analysis, which shows stable estimates across bandwidth choices, provides some reassurance against this concern.

The nonparametric RDD results add nuance to the comparison. At very narrow bandwidths (17–27 meters), the discontinuity at retained boundaries (−24.4 percent) is roughly twice that at gained boundaries (−11.5 percent). This suggests that at the sharpest spatial scale, long-standing designation does create a larger price gap. The social and informational infrastructure around a long-standing boundary—residents' mental maps, real estate agents' marketing language, municipal service boundaries, and the physical visibility of differential public investment—may create an additional layer of price segmentation that new boundaries have not yet developed.

8.3 Transaction Volume

A complementary mechanism through which designation could affect property markets is through transaction volume rather than price. If QPV designation reduces demand for housing inside the zone, one would expect fewer transactions per unit area relative to comparable areas outside. The McCrary-style density test in Figure 3 provides direct evidence on this question: transaction density varies smoothly through the boundary for both gained and retained zones, with no evidence of a discontinuous jump at the cutoff. The smooth density profile suggests that the price effect operates through the valuation of properties rather than through changes in market thickness or liquidity. Buyers and sellers transact at similar rates on both sides of the boundary; it is the equilibrium price, not the volume of activity, that adjusts to reflect designation status.

8.4 Implications for Place-Based Policy Design

The documented price differentials at QPV boundaries raise questions about the design of place-based policies, even though this study cannot determine what share of the differential is attributable to designation versus pre-existing neighborhood differences. If designation creates or reinforces a price penalty, the net welfare effect of the policy is ambiguous: investment in designated neighborhoods may be partly offset by the signal that designation sends to housing markets.

The 6–8 percent boundary price differential documented here is consistent with either a labeling effect, a reflection of underlying neighborhood quality differences, or a combination. Distinguishing between these interpretations requires pre-reform data that would allow a difference-in-discontinuities design around the 2015 boundary change. The existence of a sharp, measurable price differential at QPV boundaries nonetheless provides a useful input for future research on the equilibrium effects of place-based policy targeting.

9. Conclusion

This paper documents sharp property price differentials at France's QPV priority zone boundaries. Using 2.1 million transactions and 1,236 unique boundary segments, I find that properties inside QPV zones sell for 6–8 percent less than observationally similar properties just outside, after controlling for property characteristics and smooth spatial price gradients. The differential is similar for zones newly designated in 2015 and zones that retained priority status from the predecessor ZUS system. Nonparametric RDD estimates at very narrow bandwidths show a larger discontinuity at retained boundaries (−24.4 percent) than at gained boundaries (−11.5 percent), suggesting that the immediate boundary effect is larger for long-standing priority zones.

Two limitations shape the interpretation of these findings. First, the cross-sectional design cannot distinguish the effect of designation from pre-existing neighborhood quality differences. QPV boundaries were drawn along income gradients, ensuring that the inside was systematically poorer than the outside at the moment of designation. The estimated boundary differential therefore combines any designation effect with the capitalized value of these pre-existing differences. A credible causal design would require pre-reform transaction data to estimate a difference-in-discontinuities around the 2015 boundary change. Second, the commune-level classification of gained versus retained zones introduces measurement error, since it cannot confirm that specific QPV neighborhoods overlap with former ZUS boundaries.

Despite these limitations, the boundary discontinuity design establishes that QPV zone boundaries coincide with economically significant price discontinuities. Whether these discontinuities are caused by, or merely coincide with, the policy boundary is a question for future research with richer data. The finding that administratively drawn boundaries—even those defined by a transparent income formula—are associated with sharp price segmentation should inform the design of place-based policies that rely on public boundary-based targeting.

References

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf, "The Economics of Density: Evidence from the Berlin Wall," *Econometrica*, 2015, *83* (6), 2127–2189.

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.

Black, Sandra E., "Do Better Schools Matter? Parental Valuation of Elementary Education," *Quarterly Journal of Economics*, 1999, *114* (2), 577–599.

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

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik, "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 2014, *82* (6), 2295–2326.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz, "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment," *American Economic Review*, 2016, *106* (4), 855–902.

Chyn, Eric, "Moved to Opportunity: The Long-Run Effects of Public Housing Demolition on Children," *American Economic Review*, 2018, *108* (10), 3028–3056.

Coulson, N. Edward and Robin M. Leichenko, "Historic Preservation and the Revitalization of Neighborhoods," *Urban Studies*, 2008, *45* (8), 1621–1641.

Duguet, Emmanuel, Noam Leandri, Yannick L'Horty, and Pascale Petit, "Are Young French Jobseekers of Ethnic Immigrant Origin Discriminated Against? A Controlled Experiment in the Paris Area," *Annals of Economics and Statistics*, 2010, *99/100*, 187–215.

Galster, George C., "The Mechanism(s) of Neighbourhood Effects: Theory, Evidence, and Policy Implications," *Neighbourhood Effects Research: New Perspectives*, 2012, pp. 23–56.

Gibson, Matthew and Jamie T. Mullins, "The Effects of Flood Zone Designation on Housing Markets," *American Economic Journal: Economic Policy*, 2019.

Gobillon, Laurent, Thierry Magnac, and Harris Selod, "The Impact of Enterprise Zones on Establishment Location Decisions and Labor Market Outcomes: Evidence from France," *Journal of Urban Economics*, 2012, *72* (1), 1–17.

Hanson, Andrew and Shawn Rohlin, "Do Enterprise Zones Really Boost Employment? An Assessment Using Propensity Score Matching," *Journal of Urban Economics*, 2009, *65* (3), 237–248.

Keele, Luke J. and Rocio Titiunik, "Geographic Boundaries as Regression Discontinuities," *Political Analysis*, 2015, *23* (1), 127–155.

Lynch, Devon and Jeffrey S. Zax, "The Effects of Enterprise Zones on Property Values," *Regional Science and Urban Economics*, 2005, *35* (3), 286–306.

Mayer, Thierry and Corentin Trevien, "The Effects of Urban Renewal on Retail Activity: Evidence from the French National Program for Urban Renewal," *Journal of Urban Economics*, 2017, *100*, 51–64.

Neumark, David and Helen Simpson, "Do Place-Based Policies Matter?," *Journal of Economic Perspectives*, 2015, *29* (3), 225–252.

ONPV, "Rapport Annuel de l'Observatoire National de la Politique de la Ville," Technical Report, Observatoire National de la Politique de la Ville 2015.

Turner, Matthew A., Andrew Haughwout, and Wilbert van der Klaauw, "Land Use Regulation and Welfare," *Econometrica*, 2014, *82* (4), 1341–1403.

Appendix

A. Additional Results

Bandwidth sensitivity detail. The following table reports the full set of bandwidth sensitivity estimates underlying Figure 2.

Table A1: Bandwidth Sensitivity: Full Results

Bandwidth (m)	Group	Estimate	SE	N
200	Gained	−0.058	(0.028)	440,247
200	Retained	−0.053	(0.019)	440,247
500	Gained	−0.081	(0.031)	848,565
500	Retained	−0.059	(0.025)	848,565
1,000	Gained	−0.058	(0.028)	1,437,069
1,000	Retained	−0.048	(0.023)	1,437,069
1,500	Gained	−0.056	(0.029)	1,862,076
1,500	Retained	−0.050	(0.026)	1,862,076
2,000	Gained	−0.066	(0.030)	2,122,991
2,000	Retained	−0.052	(0.029)	2,122,991

Notes: Estimates from the preferred specification (with property controls and boundary + year FE) at different bandwidth thresholds. Standard errors clustered at boundary level.

B. Data Construction Details

The analysis panel is constructed from five annual DVF files (2020–2024), each downloaded by department from `files.data.gouv.fr/geo-dvf`. Transactions are geocoded, projected to Lambert-93 (EPSG:2154), and spatially matched to the nearest QPV zone boundary. The distance computation uses polygon boundaries cast as linestrings (`st_boundary` in the `sf` package), ensuring that distances are measured to the polygon edge rather than the polygon interior.

QPV zone classification uses commune-name matching between the QPV shapefile's `COMMUNE_QP` field and the ZUS commune list from the SIG Ville database. Multi-commune QPV zones are classified as "retained" if any listed commune matches a ZUS commune name.

Acknowledgements

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

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

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