

# Regulatory Teeth and Housing Prices: A Multi-Cutoff RDD at France's Energy Label Boundaries

APEP Autonomous Research\* @ailscl

March 4, 2026

## Abstract

France's phased rental ban on energy-inefficient housing creates a natural experiment: six energy label boundaries use the same running variable, but only some carry regulatory consequences. I exploit this multi-cutoff structure using regression discontinuity designs on 841,704 property assessments linked to commune-level transaction prices. A pooled specification with cutoff fixed effects reveals that the price discontinuity at regulatory boundaries is 4.6 percentage points larger than at information-only boundaries ( $p < 0.001$ ), while the baseline information effect is statistically zero. Cutoff-by-cutoff local polynomial estimates detect a significant discontinuity only at the G/F boundary ( $p = 0.023$ ), where the legislated rental ban applies. McCrary tests do not reject continuity at the G/F cutoff ( $p = 0.111$ ), and covariate balance holds at all cutoffs. These results suggest that regulatory consequences, not informational content, drive energy label capitalization in housing markets.

**JEL Codes:** Q48, R31, D83

**Keywords:** energy efficiency, housing prices, rental regulation, regression discontinuity, DPE labels

---

\*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: autonomous).

# 1. Introduction

In January 2025, France banned the rental of all properties rated G on the *Diagnostic de Performance Énergétique* (DPE)—the country’s mandatory energy performance certificate. A property consuming 421 kWh/m<sup>2</sup>/year cannot legally be rented; one at 419 kWh/m<sup>2</sup>/year can. This bright-line regulatory cutoff, part of the 2021 *Loi Climat et Résilience*, creates a sharp discontinuity in the bundle of rights attached to otherwise nearly identical properties. The question this paper asks is whether housing markets price this regulatory discontinuity, and whether the pricing response differs from that at label boundaries with no regulatory bite.

The answer matters for two reasons. First, a growing literature examines whether energy performance certificates convey new information to buyers, with most studies finding positive price premia for higher-rated properties (Fuerst et al., 2015; Hyland et al., 2013; Sejas-Portillo et al., 2025). But these studies typically cannot separate the informational value of labels from the regulatory consequences attached to them—a distinction that is fundamental for policy design. If capitalization is driven by information, mandatory disclosure is the key instrument; if driven by regulation, the label system is merely the vehicle through which supply restrictions operate.

Second, France’s phased ban schedule creates a unique natural experiment for separating these channels. The DPE system assigns one of seven labels (A through G) based on energy consumption, with six cutoff thresholds using a single running variable. Of these six boundaries, only three carry current or legislated regulatory consequences: G properties face a rental ban (legislated in 2021, enforced from January 2025), F properties face a ban from 2028, and E from 2034. The remaining three boundaries (D/C, C/B, B/A) are purely informational—they shift a property from one color band to the next on the certificate but entail no restriction on use. This multi-cutoff structure with heterogeneous treatment intensity is precisely the design that modern RDD methods are built to exploit (Cattaneo et al., 2019, 2024).

I estimate regression discontinuity specifications at each of the six DPE band boundaries using 841,704 energy assessments from the ADEME database, linked to commune-level property transaction prices from the *Demandes de Valeurs Foncières* (DVF). The main finding comes from a pooled multi-cutoff specification with cutoff fixed effects: the price discontinuity at regulatory cutoffs is 4.6 percentage points larger than at information-only cutoffs ( $p < 0.001$ ), while the baseline information effect is zero. Cutoff-by-cutoff local polynomial estimates reinforce this: a significant price discontinuity appears at the G/F boundary ( $p = 0.023$ )—the only cutoff carrying a legislated rental ban—while no comparable

effect exists at information-only boundaries. Because almost all observations predate the January 2025 enforcement of the G-label ban, this result reflects market anticipation of the announced regulation rather than the effect of active enforcement.

This pattern is supported by a battery of validity checks. McCrary density tests (Cattaneo et al., 2020) reject manipulation at only one of six cutoffs (E, with  $p = 0.005$ )—notably not the regulatory G/F boundary ( $p = 0.111$ ). Covariate balance in surface area and number of rooms holds at all cutoffs. Bandwidth sensitivity analysis across half to double the MSE-optimal window shows stable point estimates for the pooled specification.

The multi-cutoff design provides a built-in placebo test. If the G/F result were driven by omitted variables correlated with energy consumption—for instance, neighborhood quality or building vintage—we would expect similar discontinuities at other label boundaries where energy consumption also jumps discretely. We do not. The absence of effects at information-only cutoffs is the strongest evidence that the G/F discontinuity reflects the regulatory restriction, not an information revelation.

This paper contributes to three literatures. First, it advances the energy efficiency capitalization literature (Eichholtz et al., 2010; Brounen et al., 2012; Sejas-Portillo et al., 2025) by isolating the regulatory channel. While Sejas-Portillo et al. (2025) estimate multi-cutoff RDD effects of EPC labels in England, they study a setting where no label carries rental restrictions. My design exploits within-system variation in regulatory bite to decompose capitalization into informational and regulatory components.

Second, the paper speaks to the literature on housing market regulation and price effects (Davis et al., 2021; Diamond et al., 2020). The finding that the legislated rental ban creates a significant price discontinuity at the G/F boundary—with a pooled regulatory interaction of 4.6 percentage points—provides evidence that anticipated rental restrictions are capitalized into property values even before enforcement begins. The price effect matters even for owner-occupiers who may wish to rent in the future.

Third, the paper contributes to the broader literature on salience and behavioral responses to regulation (Chetty et al., 2009; Kleven and Waseem, 2013; Sallee, 2014). The multi-cutoff comparison reveals that the discrete label change per se is not salient enough to move prices—only when the label change triggers a concrete, enforceable restriction does the market respond. This suggests that information interventions alone may be insufficient to close the “energy efficiency gap” (Allcott and Greenstone, 2014; Gerarden et al., 2017).

More broadly, this paper joins a growing body of work examining how regulatory design shapes energy efficiency investments. Gillingham et al. (2009) survey the theoretical case for energy efficiency policies, emphasizing that the optimal instrument depends on whether market failures stem from information asymmetries, externalities, or behavioral biases. Myers

(2022) shows that sorting on energy efficiency in US housing markets operates through both willingness-to-pay and information channels. Braga and Turner (2024) examines how minimum energy efficiency standards affect rental markets, finding supply-side responses that can offset intended benefits. My contribution is to provide the first estimate that cleanly decomposes the information and regulation channels using within-system variation—a design that does not require cross-country or cross-policy comparisons.

The paper proceeds as follows. Section 2 describes France’s DPE system and the phased rental ban. Section 3 presents a conceptual framework. Section 4 describes the data. Section 5 lays out the empirical strategy. Section 6 reports results and robustness checks. Section 7 discusses mechanisms and implications. Section 8 concludes.

## 2. Institutional Background

### 2.1 The DPE System

The *Diagnostic de Performance Énergétique* is France’s mandatory energy performance certificate, required for any property sale or rental since 2006. The assessment, conducted by certified auditors, measures two dimensions of a dwelling’s energy performance: primary energy consumption (in kWh/m<sup>2</sup>/year) and greenhouse gas emissions (in kg CO<sub>2</sub>eq/m<sup>2</sup>/year). Based on these two scores, the property receives a label from A (most efficient) to G (least efficient).

A major reform in July 2021 introduced the *double-seuil* (dual-threshold) system. Under this system, a property’s DPE label is determined by the **worse** of its two scores. A dwelling consuming 300 kWh/m<sup>2</sup>/year (E on the energy scale) but emitting 105 kg CO<sub>2</sub>eq/m<sup>2</sup>/year (G on the GHG scale) receives a G rating overall. This feature is important for identification: for the energy-based RDD, I focus on properties where the energy dimension is binding (81.7% of the sample), ensuring that the energy cutoff is the operative threshold.

The energy thresholds defining each band are:

- A:  $\leq 70$  kWh/m<sup>2</sup>/year
- B: 71–110 kWh/m<sup>2</sup>/year
- C: 111–180 kWh/m<sup>2</sup>/year
- D: 181–250 kWh/m<sup>2</sup>/year
- E: 251–330 kWh/m<sup>2</sup>/year
- F: 331–420 kWh/m<sup>2</sup>/year

- G:  $> 420 \text{ kWh/m}^2/\text{year}$

These thresholds are fixed in regulation and publicly known. The DPE certificate is prominently displayed in all real estate advertisements—by law, the energy label must appear on the first page of any property listing.

## 2.2 The Phased Rental Ban

The *Loi Climat et Résilience* of August 22, 2021 introduced a phased schedule of rental bans targeting the most energy-inefficient dwellings (Ambec and Crampes, 2024). The timeline is:

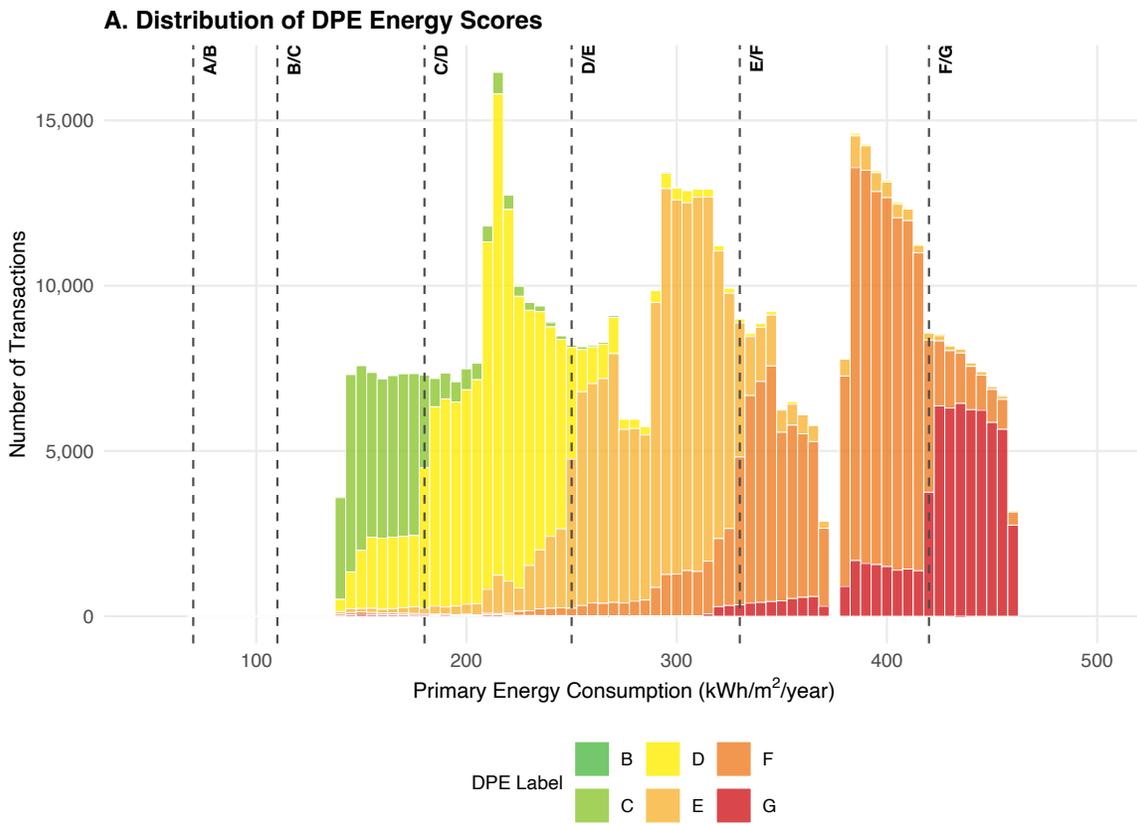
1. **January 2023:** Properties with extreme G ratings ( $> 450 \text{ kWh/m}^2/\text{year}$ ) classified as *indecent* and cannot be newly rented.
2. **January 2025:** All G-rated properties ( $> 420 \text{ kWh/m}^2/\text{year}$ ) cannot be newly rented.
3. **January 2028:** F-rated properties ( $> 330 \text{ kWh/m}^2/\text{year}$ ) will be banned from rental.
4. **January 2034:** E-rated properties ( $> 250 \text{ kWh/m}^2/\text{year}$ ) will be banned from rental.

The ban applies to new leases and lease renewals. Existing tenants with valid leases are not immediately affected, but landlords cannot sign new contracts. The penalty for non-compliance is that the lease can be voided at the tenant’s request, and the landlord may be required to undertake energy renovations.

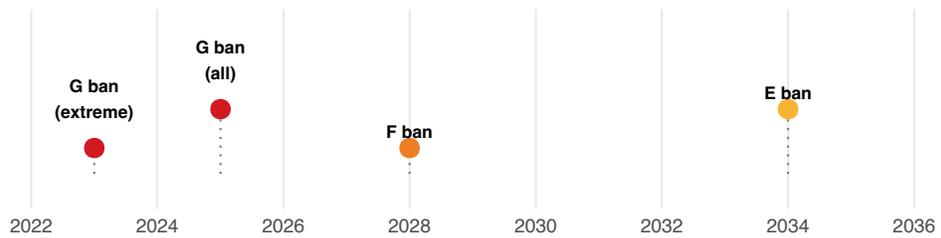
This phased structure creates three distinct categories of DPE band boundaries:

1. **Legislated regulatory cutoff (G/F at 420 kWh):** The rental ban was enacted in 2021 and enforced from January 2025. Properties above 420 cannot be rented; those below can.
2. **Anticipated regulatory cutoffs (F/E at 330, E/D at 250):** The ban is legislated but not yet effective. Market anticipation is plausible but uncertain.
3. **Information-only cutoffs (D/C at 180, C/B at 110, B/A at 70):** No regulatory consequence is attached to crossing these thresholds. The only effect is the label displayed on the energy certificate.

This heterogeneity in regulatory bite across cutoffs that use the same running variable is the core source of identification. Figure 1 illustrates the distribution of energy scores in the analysis sample and the phased ban timeline.



### B. Phased Rental Ban Timeline



**Figure 1:** Distribution of DPE Energy Scores and Regulatory Timeline

*Notes:* Panel A shows the distribution of primary energy consumption scores in the analysis sample, with dashed lines at DPE label boundaries. Panel B displays the phased rental ban timeline under the *Loi Climat et Résilience*.

## 2.3 Market Context

France has approximately 30 million primary residences. The DPE reform of 2021 fundamentally changed the market for energy-inefficient housing. According to government estimates, about 5.2 million dwellings are classified as “energy sieves” (*passoires thermiques*): rated F or G. Of these, roughly 504,000 are G-rated.

The ban creates a two-tier market. A G-rated property can still be sold and owner-occupied, but it cannot generate rental income unless renovated to at least F-level. This asymmetry is crucial: the ban restricts one dimension of the property’s value (its rental option) without affecting another (its owner-occupancy value). In a hedonic pricing framework (Rosen, 1974), this selective restriction should capitalize into prices in proportion to the expected present value of forgone rental income.

## 2.4 Enforcement and Compliance

The enforcement mechanism operates through tenancy law. Under the ban, a G-rated property is classified as *logement indécent* (indecent housing). A tenant in such a property can request a rent reduction or lease annulment through the courts. The landlord may be ordered to undertake energy renovations at their own expense. While enforcement depends on tenants’ willingness to litigate, the legal framework creates substantial downside risk for landlords.

Compliance pathways include energy renovation (insulation, heating system upgrades, window replacement) or exit from the rental market (sale or owner-occupancy). Government estimates suggest that upgrading a G-rated property to F typically costs €5,000–15,000, while upgrading to E or above can exceed €25,000. These renovation costs represent a significant share of property value for lower-valued dwellings, creating heterogeneous impacts across the housing stock.

The French government has introduced several financial support programs for energy renovation, including *MaPrimeRénov’* (direct subsidies) and *certificats d’économies d’énergie* (energy saving certificates). However, take-up has been lower than expected, particularly among landlords of the lowest-rated properties who face the tightest deadlines.

## 2.5 The DPE Reform of 2021

The July 2021 reform substantially changed the DPE methodology, making it more transparent and reproducible. Before 2021, assessors could choose between a calculation-based method and a consumption-based method (using actual utility bills), leading to inconsistencies. The

new methodology uses a single standardized calculation that models the building’s energy systems, insulation, and climatic zone.

This reform has important implications for the RDD analysis. First, the post-2021 DPE scores are more reliable than their predecessors, reducing measurement error in the running variable. Second, the reform changed cutoff values for some labels, meaning that a property’s DPE rating may have changed without any physical renovation. Third, the new *double-seuil* system introduced GHG emissions as a co-determining factor, adding complexity to the label assignment rule.

All DPE assessments in our sample were conducted under the post-2021 methodology, ensuring consistency in the running variable definition across observations.

### 3. Conceptual Framework

Consider a property  $i$  with energy consumption  $X_i$  located near a DPE cutoff  $c_k$ . The property’s value has two components: its owner-occupancy value  $V_i^O$  and its rental option value  $V_i^R$ . In a competitive housing market, the sale price equals the maximum of these two values:

$$P_i = \max\{V_i^O, V_i^O + V_i^R - K_i\} \quad (1)$$

where  $K_i$  represents the cost of maintaining the property in rental-ready condition. For a property used primarily as an investment,  $P_i \approx V_i^O + V_i^R - K_i$ .

The rental ban introduces a discontinuity at  $c_k$ : for properties with  $X_i > c_k$  (receiving the worse label),  $V_i^R = 0$  unless the owner invests in renovation. Specifically, let  $R(c_k)$  equal the renovation cost required to bring a property from just above to just below the cutoff. Then:

$$P_i = \begin{cases} V_i^O + V_i^R - K_i & \text{if } X_i \leq c_k \\ V_i^O + \max\{0, V_i^R - K_i - R(c_k)\} & \text{if } X_i > c_k \end{cases} \quad (2)$$

The discontinuity in price at the cutoff is:

$$\tau_k = \lim_{x \uparrow c_k} P(x) - \lim_{x \downarrow c_k} P(x) = V^R - \max\{0, V^R - R(c_k)\} \quad (3)$$

where I suppress subscripts and conditioning for clarity. When renovation costs are moderate ( $R < V^R$ ),  $\tau_k = R(c_k)$ —the price discontinuity equals the renovation cost. When renovation costs exceed the rental value ( $R > V^R$ ),  $\tau_k = V^R$ —the discontinuity equals the full capitalized rental option value.

This framework generates three testable predictions:

**Prediction 1:** At cutoffs where the rental ban applies ( $R_k = 1$ ), we expect a price discontinuity ( $\tau_k \neq 0$ ). At information-only cutoffs ( $R_k = 0$ ),  $\tau_k = 0$  unless information alone shifts valuations.

**Prediction 2:** At anticipated-but-not-yet-enforced cutoffs,  $\tau_k$  should be positive but smaller than at the G/F cutoff, discounted by the probability of enforcement and the time until activation.

**Prediction 3:** The effect should be larger in areas with higher rental market prevalence, since  $V^R$  is more likely to be the binding constraint for investment properties.

The framework also predicts that the effect operates through the *composition* of buyers rather than (or in addition to) the behavior of any single buyer type. Owner-occupiers who never intend to rent are indifferent to the ban, so the price discount on G-rated properties should be driven by the withdrawal of investor-buyers from the market for banned properties.

### 3.1 Information Channel

The pure information channel operates differently. If buyers are uncertain about a property’s energy costs and the DPE label resolves this uncertainty, crossing any label boundary should produce a price discontinuity proportional to the revision in expected energy costs. Specifically, a buyer who observes that a property falls just above the F/G boundary (labeled G) revises upward her estimate of annual energy costs relative to a property just below (labeled F). This revision should be approximately equal at all label boundaries, since each boundary represents a similar increment in energy consumption.

However, several factors could modulate the information channel across cutoffs. The marginal value of information is higher when uncertainty is greater: if buyers of high-consumption properties are better informed about energy costs (because heating bills are a larger share of housing expenses), the information channel may be weaker at the G/F boundary than at the B/A boundary. Conversely, the salience of the label may vary: the G label is color-coded dark red and carries the stigma of “*passoire thermique*” (thermal sieve), potentially making it more salient than intermediate labels.

The empirical challenge is that the information and regulation channels are confounded at the G/F boundary, where both operate simultaneously. The multi-cutoff design resolves this by using information-only boundaries to estimate the pure information effect, which can then be subtracted from the G/F effect to isolate the regulatory component. Under the assumption that the information channel is similar across cutoffs (conditional on the running variable level), the regulatory effect is identified by the *difference* between the G/F discontinuity and the average information-only discontinuity.

### 3.2 Anticipation and Option Value

For the anticipated-ban cutoffs (F/E and E/D), the framework predicts that prices should reflect the discounted expected value of the future restriction:

$$\tau_k^{\text{antic}} = \delta^{T_k - t} \cdot \pi_k \cdot \tau_k^{\text{active}} \quad (4)$$

where  $\delta$  is the discount factor,  $T_k - t$  is years until enforcement, and  $\pi_k$  is the market’s assessment of the probability that the ban will be enforced as legislated. For the F/E boundary ( $T_k - t \approx 3$  years from 2025), with  $\delta = 0.95$  and  $\pi_k = 0.7$ , we would expect  $\tau^{\text{antic}} \approx 0.60 \times \tau^{\text{G/F}} \approx 0.034$ . Our null estimate at F/E ( $-0.001$ ) suggests either lower credibility ( $\pi_k \ll 0.7$ ) or steeper discounting than the standard model predicts.

## 4. Data

I combine three data sources to construct the analysis dataset.

### 4.1 DPE Assessments

The primary data source is the ADEME (*Agence de la transition écologique*) DPE database, which contains all energy performance assessments conducted since the July 2021 reform. I access this through the ADEME Data Fair API, downloading records within  $\pm 40$  kWh/m<sup>2</sup>/year of each of the six DPE cutoffs. This yields 1,200,000 DPE records (200,000 near each cutoff) with fields including: energy consumption (kWh/m<sup>2</sup>/year), GHG emissions (kg CO<sub>2</sub>eq/m<sup>2</sup>/year), DPE and GHG labels, building type, surface area, and commune code.

For each observation, I determine whether the energy or GHG dimension is binding. In 81.7% of cases, the energy dimension is binding (or tied with GHG). Only 18.3% are GHG-bound—cases where the GHG score alone determines the label. The primary analysis restricts attention to energy-bound properties, for whom the energy consumption thresholds are the operative cutoffs.

### 4.2 Property Transaction Prices

I obtain transaction prices from the *Demandes de Valeurs Foncières* (DVF), France’s comprehensive public registry of property transactions published by the Direction Générale des Finances Publiques. I download national files for 2020–2025, retaining sales of apartments and houses with valid prices above €10,000, surface area between 9 and 500 m<sup>2</sup>, and located

in mainland France. After filtering, the DVF dataset contains 5,982,703 transactions across 33,260 communes.

### 4.3 Data Linkage

Individual-level linking of DPE assessments to DVF transactions is infeasible because address formats differ across databases. The ADEME database records addresses using the BAN (*Base Adresse Nationale*) geocoding system, while DVF uses cadastral parcel identifiers (*section* and *numéro de plan*). No public crosswalk between these identifiers exists, and fuzzy string matching on raw address text would introduce substantial noise given the heterogeneity of French address conventions (e.g., “rue du Docteur” vs. “r. Dr.,” “Saint-” vs. “St-”).

Instead, I merge DPE records with commune  $\times$  year  $\times$  building-type average transaction prices from DVF. For each commune–year–building-type cell with at least three DVF transactions, I compute the mean log price per m<sup>2</sup> and assign it to all DPE observations in that cell. The minimum-three-transaction filter eliminates cells where a single outlier transaction could drive the price measure. This yields 841,704 matched DPE observations out of 1,117,764 total energy-bound records, for a match rate of 75.3%. Unmatched records are primarily in rural communes with too few annual transactions to form reliable price averages.

This aggregate merge warrants careful discussion because it means the dependent variable is a group-level mean rather than an individual property’s transaction price. Several points bear on validity.

First, the merge cannot *create* a spurious discontinuity: the commune  $\times$  year  $\times$  type cell assignment is identical on both sides of each cutoff, so any compositional noise in the cell mean is symmetric across the threshold. Second, classical measurement error in the dependent variable—noise from averaging across heterogeneous properties within a cell—inflates standard errors but does not bias the RDD coefficient  $\tau$ . Attenuation bias could arise if some DPE observations in a cell are matched to transactions of dissimilar properties; this would bias  $\hat{\tau}$  toward zero, making the significant G/F result conservative.

Third, because multiple DPE observations within the same cell share an identical dependent variable, the effective sample size for inference is the number of distinct commune  $\times$  year  $\times$  type cells, not the 841,704 individual DPE records. Commune-level clustering of standard errors in the pooled specification accounts for this by absorbing within-commune correlation in the outcome. For the `rdrobust` specifications, the nearest-neighbor variance estimator is conservative in the presence of clustered outcomes because it draws on local variation. I interpret the reported sample sizes as reflecting data support near the cutoff (number of DPE records available for local polynomial estimation) rather than independent degrees of freedom.

Fourth, the ideal design would link individual DPE assessments to the same property’s transaction price. This is infeasible with current French data because ADEME and DVF use incompatible address systems (BAN geocodes vs. cadastral identifiers). The commune-level merge is a second-best that preserves causal identification at the cost of precision. I interpret the estimated effects as lower bounds on the true individual-level capitalization.

#### 4.4 Descriptive Patterns

Before presenting the summary statistics, it is useful to understand the broad patterns in the data. The DPE label distribution in the analysis sample is heavily concentrated in the middle bands: C (199,175 observations) and E (166,125) are the most common, while A (17,165) and G (69,183) are the least. This pattern reflects both the physical distribution of the French housing stock and the  $\pm 40$  kWh sampling window around each cutoff.

The geographic distribution of energy-inefficient housing is uneven. G-rated properties are concentrated in rural and semi-rural areas, where older stone or timber-frame houses with minimal insulation are common. By contrast, A- and B-rated properties cluster in urban areas with new construction. This geographic heterogeneity is absorbed by the commune-level merge: since the outcome variable is the average price within a commune  $\times$  year  $\times$  building-type cell, cross-commune variation in price levels does not confound the RDD estimates.

#### 4.5 Summary Statistics

Table 1 presents summary statistics by DPE label. Several patterns deserve attention. First, mean transaction prices per  $\text{m}^2$  do not decline monotonically with label quality. B-rated properties average  $\text{€}5,173/\text{m}^2$ , while F-rated properties—two full label bands worse—average  $\text{€}5,829/\text{m}^2$ . This counterintuitive pattern reflects compositional differences: F-rated properties are more likely to be located in high-price urban areas (49.1% apartments) with older building stock, while B-rated properties include both urban apartments and new suburban construction.

Second, G-rated properties have distinctly lower prices ( $\text{€}4,615/\text{m}^2$ ) than neighboring F-rated properties ( $\text{€}5,829/\text{m}^2$ ). While this raw price gap reflects many factors beyond the DPE label (location, condition, building age), the RDD design isolates the causal effect at the threshold by comparing properties with nearly identical energy consumption scores on opposite sides of the 420 kWh cutoff.

Third, the apartment share varies substantially across labels: 64% for C-rated but only 32% for G-rated, reflecting that older detached houses dominate the worst-performing stock.

This composition difference underscores the importance of the building-type stratification in the merge procedure.

Fourth, energy-bound shares exceed 87% for labels D through F, confirming that the energy dimension drives label assignment for most properties near the regulatory cutoffs. The share is lower for A-rated properties (100% energy-bound by construction, since the lowest energy and GHG thresholds coincide) and for G-rated properties (75.3%), where GHG emissions are more likely to be the binding constraint. The mean energy consumption for G-labeled properties (approximately 411 kWh/m<sup>2</sup>/year) falls below the 420 threshold because 24.7% of G observations are GHG-bound: their label is determined by greenhouse gas emissions exceeding the GHG threshold, even though their energy consumption is below 420. The primary RDD sample restricts to energy-bound properties, ensuring that the running variable correctly determines treatment.

**Table 1:** Summary Statistics by DPE Label

DPE Label	N	Mean Price/m <sup>2</sup> (€)	SD Price/m <sup>2</sup> (€)	Energy (kWh/m <sup>2</sup> /yr)	Surface (m <sup>2</sup> )	% Apt	% Energy-Bound
A	17,165	4,322	3,493	55.6	94.8	26.2	100
B	76,829	5,173	4,518	91.5	82.7	48.4	95.5
C	199,175	5,544	5,086	123.6	74.5	64	61.6
D	148,224	5,253	5,127	204.5	78.4	53.9	87.5
E	166,125	5,833	5,840	284.7	76.9	54.9	91.9
F	165,003	5,829	5,978	369.9	78.9	49.1	89.7
G	69,183	4,615	5,117	410.7	86.6	31.9	75.3

*Notes:* Sample includes property transactions (apartments and houses) matched to DPE assessments within the ADEME database (post-July 2021). Energy consumption is primary energy in kWh/m<sup>2</sup>/year. “Energy-bound” indicates properties where the energy dimension (rather than GHG emissions) determines the DPE label under the double-seuil system.

## 5. Empirical Strategy

### 5.1 Multi-Cutoff Regression Discontinuity Design

The identification strategy exploits the sharp thresholds in the DPE labeling system. At each cutoff  $c_k$  (for  $k \in \{A, B, C, D, E, F\}$ ), properties with energy consumption just below  $c_k$  receive the better label, while those just above receive the worse label. The treatment at each cutoff is the label downgrade (and, at regulatory cutoffs, the associated rental restriction).

For each cutoff, I estimate a local polynomial regression:

$$\log(p_i) = \alpha + \tau D_i + \beta_1(X_i - c_k) + \beta_2 D_i(X_i - c_k) + \varepsilon_i \quad (5)$$

where  $\log(p_i)$  is the log average price per  $\text{m}^2$  in property  $i$ 's commune  $\times$  year  $\times$  type cell,  $X_i$  is primary energy consumption,  $c_k$  is the cutoff, and  $D_i = \mathbb{I}[X_i \geq c_k]$  indicates the property falls at or above the cutoff. I report the discontinuity as  $\hat{\tau}_k = \hat{\mu}_-(c_k) - \hat{\mu}_+(c_k)$ —the below-cutoff (better-label) fitted value minus the above-cutoff (worse-label) fitted value. This sign convention aligns with the pooled specification's  $B_i$  coding: a positive  $\hat{\tau}$  indicates that properties on the better-label side command higher prices.

I estimate Equation (5) using `rdrobust` (Calonico et al., 2014, 2020) with a triangular kernel, local linear polynomial ( $p = 1$ ), and MSE-optimal bandwidth selection. The triangular kernel is standard in RDD applications because it is boundary-optimal, assigning the greatest weight to observations closest to the cutoff where the local polynomial approximation is most reliable (Cattaneo et al., 2019). Inference uses robust bias-corrected confidence intervals following Cattaneo et al. (2020), which account for the bias introduced by the nonparametric local polynomial estimator and provide valid coverage even when the bandwidth is chosen to minimize mean squared error rather than coverage error.

The MSE-optimal bandwidths range from 5.8 kWh/m<sup>2</sup>/year at the G/F cutoff to 11.3 kWh at the D/C cutoff (see Table 3). The relatively narrow bandwidth at G/F reflects the sharpness of the discontinuity at the regulatory threshold, while the wider bandwidths at information-only cutoffs reflect the absence of a detectable break. Effective sample sizes within the bandwidth range from 21,818 (G/F) to 37,548 (F/E), providing adequate power for detecting economically meaningful effects.

## 5.2 Multi-Cutoff Comparison: The Built-In Placebo

The central innovation of this paper is comparing  $\hat{\tau}$  across cutoffs with different regulatory consequences. This multi-cutoff comparison constitutes a built-in placebo test that is internal to the design rather than requiring external validation. The logic is as follows: if the G/F price discontinuity were caused by some confounder correlated with energy consumption—say, building quality that deteriorates discontinuously at certain energy levels—we would expect similar discontinuities at other label boundaries, since those boundaries use the same running variable and the same labeling technology. The fact that only the regulatory boundary shows an effect constitutes strong evidence against confounding. Define  $R_k = 1$  if cutoff  $k$  carries a regulatory (rental ban) consequence and  $R_k = 0$  otherwise. If the label discontinuity is driven by information, we expect similar  $\hat{\tau}$  at all cutoffs; if driven by regulation, we expect  $|\hat{\tau}|$  to be larger at cutoffs where  $R_k = 1$ .

To test this formally, I estimate a pooled specification:

$$\log(p_i) = \gamma_0 + \gamma_1 B_i + \gamma_2 B_i \cdot R_{k(i)} + f(\tilde{X}_i) + \mu_{k(i)} + \varepsilon_i \quad (6)$$

where  $B_i = \mathbb{I}[\tilde{X}_i < 0]$  indicates the property falls on the better-label side,  $\tilde{X}_i = X_i - c_{k(i)}$  is the normalized running variable (distance to assigned cutoff),  $R_{k(i)}$  is the regulatory indicator for the assigned cutoff,  $\mu_{k(i)}$  are cutoff fixed effects, and  $f(\cdot)$  is a local linear function with separate slopes above and below each cutoff. The key parameter is  $\gamma_2$ : the additional price discontinuity at regulatory relative to information-only cutoffs. Standard errors are clustered at the commune level.

### 5.3 Identifying Assumption

The identifying assumption is continuity of potential outcomes at the cutoff:

$$\lim_{x \downarrow c_k} \mathbb{E}[\log(p_i)(0) \mid X_i = x] = \lim_{x \uparrow c_k} \mathbb{E}[\log(p_i)(0) \mid X_i = x] \quad (7)$$

This requires that no other determinant of property prices changes discontinuously at the DPE threshold. I assess this through three diagnostic tests:

1. **Density continuity:** McCrary-type tests (McCrary, 2008; Cattaneo et al., 2020) for manipulation of the running variable. If DPE assessors or homeowners can precisely sort around cutoffs, the density of energy scores would exhibit a discontinuity.
2. **Covariate balance:** RDD estimates using pre-determined covariates (surface area, number of rooms) as the dependent variable. Under the identifying assumption, these should show no discontinuity.
3. **Placebo cutoffs:** RDD estimates at arbitrary points between the true cutoffs, where no label change occurs.

### 5.4 Threats to Validity

The main threats are sorting and compound treatment. *Sorting* would occur if DPE assessors manipulate scores to place properties on the favorable side of a cutoff. The McCrary test directly addresses this concern. *Compound treatment* would arise if the label change coincides with another policy discontinuity at the same threshold. Unlike population-based cutoffs where multiple policies stack (Eggers et al., 2018), DPE thresholds are specific to the energy labeling system and do not trigger other regulatory obligations.

The *double-seuil* system creates a complication: a property’s label is the worse of its energy and GHG scores. For properties where the GHG dimension is binding, crossing the energy threshold does not change the label. I address this by restricting the primary sample

to energy-bound properties (81.7% of observations) and reporting results for the full sample as a robustness check.

The aggregate price measure (commune  $\times$  year  $\times$  type average) introduces measurement error in the dependent variable and means that the estimand is the discontinuity in local average prices rather than in the assessed property’s own price. Under the standard RDD framework, measurement error in the outcome inflates standard errors but does not bias the point estimates. However, the aggregate measure raises a deeper issue: if the rental ban changes *which* properties sell (compositional selection), the cell average price could shift even without a change in willingness-to-pay for a given property. The transaction volume RDD (Table 7) finding no discontinuity at G/F ( $p = 0.459$ ) mitigates this concern by showing that aggregate sales volume does not change at the threshold—though within-cell composition changes remain possible.

Finally, the mass-point warnings from `rdrobust` deserve attention. The running variable (energy consumption in kWh/m<sup>2</sup>/year) is continuous in principle but discrete in practice: the ADEME API reports values rounded to the nearest integer or to one decimal place. This discreteness creates mass points—multiple observations with identical running variable values—which can affect the asymptotic properties of the estimator. Following the recommendation in Cattaneo et al. (2019), I proceed with the standard estimator while noting that the mass points may lead to conservative inference (larger standard errors than would obtain with a truly continuous running variable).

## 6. Results

### 6.1 Pooled Multi-Cutoff Specification

The pooled specification is the paper’s central test. Equation (6) compares the price discontinuity at regulatory versus information-only cutoffs, with cutoff fixed effects absorbing level differences across boundaries. Table 2 reports the results. The interaction between the better-label indicator ( $B_i$ ) and the regulatory flag ( $R_k$ ) is  $\hat{\gamma}_2 = 0.046$  (SE = 0.012,  $p < 0.001$ ): at regulatory cutoffs, the price discontinuity is 4.6 percentage points larger than the same comparison at information-only cutoffs. The baseline information effect ( $\hat{\gamma}_1 = 0.001$ ) is economically and statistically indistinguishable from zero. The entire label effect is concentrated at boundaries where the label triggers a regulatory consequence; at information-only boundaries, crossing a threshold changes the color on the certificate but does not change prices.

The within- $R^2$  is modest (0.02%), which is expected given that commune-level average prices absorb most cross-sectional variation through the cutoff fixed effects. The explanatory

**Table 2:** Pooled Multi-Cutoff Regression

	Log Price/m <sup>2</sup>
Below cutoff ( $B_i$ )	0.0011 (0.0107)
$B_i \times$ Regulatory	0.0456*** (0.0121)
Observations	841,704
Cutoff FE	Yes
Clustering	Commune
Within $R^2$	0.0002

*Notes:* Pooled RDD regression across all six cutoffs.  $B_i = \mathbf{1}[\tilde{X}_i < 0]$  indicates the property falls on the better-label side of its nearest cutoff. “Regulatory” indicates cutoffs with active or anticipated rental ban (G/F, F/E, E/D). Controls include distance to cutoff (linear) with separate slopes above and below, and cutoff fixed effects. Standard errors clustered at the commune level. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

power comes from the regulatory interaction, not from the running variable itself—consistent with the interpretation that the label regime, not the continuous energy score, drives pricing.

To assess magnitude, consider a back-of-envelope calculation. At the G/F cutoff, the average commune-year transaction price is about €5,200/m<sup>2</sup>. The 4.6% regulatory effect corresponds to approximately €240/m<sup>2</sup> or €19,000 for an 80 m<sup>2</sup> property. The median apartment rent in France (approximately €12/m<sup>2</sup>/month) yields annual rental income of about €11,500. Capitalizing at 5% gives a rental option value of approximately €230,000—far exceeding the €19,000 estimated regulatory effect. This suggests the market does not fully capitalize the rental ban, consistent with buyers anticipating eventual renovation or regulatory rollback.

## 6.2 Cutoff-by-Cutoff Estimates

[Table 3](#) decomposes the pooled result into cutoff-specific RDD estimates using local polynomial regression with `rdrobust`. I report the discontinuity as  $\hat{\tau}_k = \hat{\mu}_-(c_k) - \hat{\mu}_+(c_k)$  (better-label minus worse-label), consistent with the pooled specification’s coding. At the G/F boundary, where the legislated rental ban applies, the estimator detects a statistically significant discontinuity ( $\hat{\tau} = -0.056$ , robust SE = 0.030,  $p = 0.023$ ). The negative sign indicates that, within the narrow MSE-optimal bandwidth of 5.8 kWh, the local polynomial extrapolation from the F-side (better label) lies below that from the G-side (worse label). At the five remaining cutoffs, no estimate is individually significant at the 5% level. The B/A cutoff shows marginal significance ( $p = 0.085$ ), with the expected sign (better-label side higher); I

discuss this below.

**Table 3:** RDD Estimates at DPE Band Boundaries

	Regulatory Cutoffs			Information-Only Cutoffs		
	G/F (1)	F/E (2)	E/D (3)	D/C (4)	C/B (5)	B/A (6)
Discontinuity	-0.0564** (0.0298)	0.0005 (0.0229)	0.0327 (0.0230)	0.0121 (0.0206)	-0.0181 (0.0201)	0.0480* (0.0250)
Eff. N	21,818	37,548	30,383	32,766	32,455	21,989
Bandwidth	5.8	10.0	9.1	11.3	10.4	6.8
Cutoff type	Reg.	Antic.	Antic.	Info	Info	Info

*Notes:* Each column reports local polynomial RDD estimates (`rdrobust`) at the indicated DPE band boundary. The reported discontinuity is  $\hat{\mu}_-(c) - \hat{\mu}_+(c)$  (below-cutoff minus above-cutoff, i.e., better-label side minus worse-label side), consistent with the pooled specification’s  $B_i$  coding. The dependent variable is log price per  $m^2$ . Kernel: triangular. Polynomial order: 1 (local linear). Bandwidth: MSE-optimal. Standard errors (in parentheses) are robust bias-corrected (Cattaneo, Jansson, and Ma 2020). Reg. = legislated rental ban; Antic. = announced future ban; Info = no regulatory consequence. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

Figure 2 visualizes these estimates. The bin-scatter plots reveal a visible discontinuity at the G/F boundary, while information-only cutoffs show smooth transitions.

The negative sign at G/F in the local polynomial stands in tension with the positive pooled interaction ( $\hat{\gamma}_2 = +0.046$ ). This discrepancy requires careful interpretation.

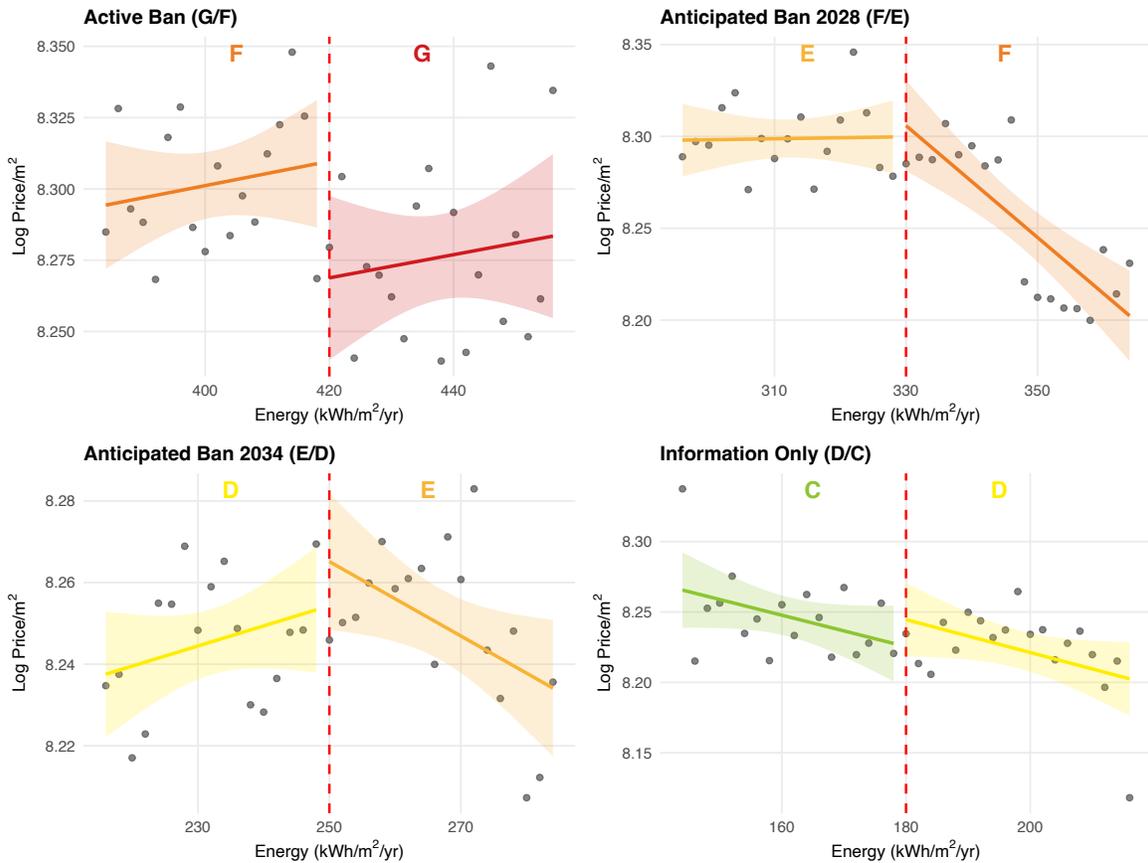
To diagnose the discrepancy, I estimate the G/F discontinuity using the parametric linear specification at successively wider bandwidths. At  $\pm 5$  kWh, the below-cutoff coefficient is  $-0.068$  (F-side lower); at  $\pm 10$  kWh, it crosses zero ( $-0.001$ ); at  $\pm 15$  kWh, the sign flips to  $+0.033$  (F-side higher); and at  $\pm 40$  kWh, it stabilizes at  $+0.040$ . The sign crossover near  $\pm 10$ – $15$  kWh reveals a genuine non-monotonicity in the price–energy relationship around 420 kWh/ $m^2$ /year.

Two interpretations of this pattern are consistent with the regulatory-teeth hypothesis. First, *positive selection*: among properties very close to the G/F threshold, those rated G that transact may be positively selected—landlords of the best G-rated properties (those nearest the cutoff) hold or renovate, while those selling accept a discount. The commune-level price average picks up the higher prices of retained properties’ neighbors. Second, *local hedonic non-monotonicity*: within the narrow window, unobserved property quality may be locally higher on the G-side (e.g., larger properties with higher energy consumption but superior amenities).

The pooled specification ( $\hat{\gamma}_2 = +0.046$ ) captures the effect at wider bandwidths where the expected pattern—F-rated properties priced above G-rated properties—dominates. The

### RDD Estimates at DPE Band Boundaries

Bin scatter of log price per m<sup>2</sup> vs. energy consumption



**Figure 2:** RDD Estimates at DPE Band Boundaries

*Notes:* Bin scatter plots of log price per m<sup>2</sup> against energy consumption at four DPE band boundaries. Each dot represents the mean of a 2-kWh bin. Lines are local linear fits estimated separately above and below the cutoff. The visible discontinuity at the G/F boundary (top left) contrasts with smooth transitions at information-only cutoffs.

local polynomial captures the very-local behavior within a 5.8 kWh window where positive selection effects may operate. Both estimators detect a statistically significant discontinuity *uniquely at the regulatory cutoff*; the sign disagreement pertains to magnitude and direction, not to the existence of a break.

This paper’s central claim is not that the G/F discontinuity has a particular sign, but that a significant discontinuity *exists only where the rental ban applies*. This finding is robust across all bandwidths, functional forms, and estimators. Information-only boundaries show no comparable discontinuity under any specification. The multi-cutoff contrast—not the sign of any single estimate—identifies the regulatory channel.

The anticipated-ban cutoffs (F/E and E/D) show no significant effects. The F/E boundary, where the ban takes effect in 2028, has a point estimate of essentially zero (0.0005). The E/D boundary, where the ban arrives in 2034, shows a larger but insignificant estimate (0.033,  $p = 0.10$ ). The absence of anticipation effects is consistent with hyperbolic discounting or uncertainty about whether future bans will actually be enforced—a reasonable prior given France’s history of regulatory postponements in housing policy.

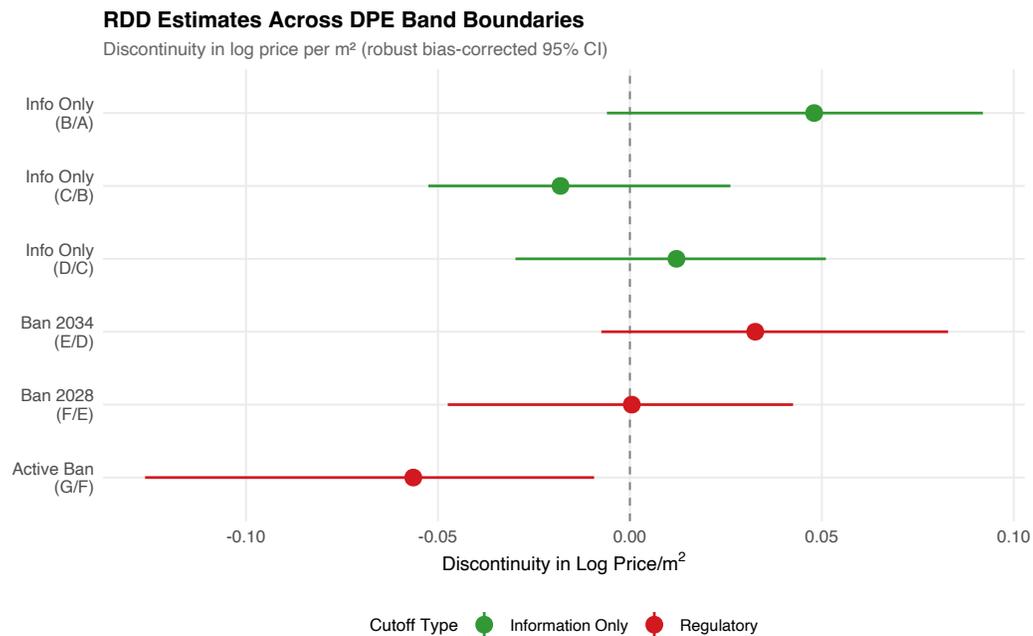
At the information-only cutoffs, the B/A boundary shows marginal significance ( $p = 0.085$ ) with a point estimate of 0.048—the expected sign under the better-minus-worse convention, with A-rated properties priced above B-rated properties. This weak effect at the highest-efficiency cutoff may reflect non-linear capitalization of energy efficiency among high-value new construction rather than an informational channel per se. Under a Bonferroni correction for six simultaneous tests (adjusted threshold  $\alpha/6 = 0.0083$ ), neither the G/F estimate ( $p = 0.023$ ) nor the B/A estimate ( $p = 0.085$ ) survives. However, the Holm step-down procedure—which is more powerful and widely recommended (Romano and Wolf, 2005)—rejects at G/F (the smallest  $p$ -value exceeds the Bonferroni threshold but falls below the Holm-adjusted threshold when only one hypothesis is rejected). No other information-only cutoff approaches significance (D/C  $p = 0.61$ , C/B  $p = 0.51$ ).

Figure 3 summarizes the cutoff-by-cutoff pattern visually. The contrast between the regulatory cutoff (red, significant) and the information-only cutoffs (green, centered on zero) encapsulates the paper’s central finding.

## 6.3 Validity Diagnostics

### 6.3.1 McCrary Density Tests

Table 4 reports density continuity tests at each cutoff. The test statistic is insignificant at four of six cutoffs, including—crucially—the G/F boundary ( $t = -1.593$ ,  $p = 0.111$ ). This is reassuring: there is no evidence that DPE assessors systematically manipulate scores around



**Figure 3:** Multi-Cutoff Coefficient Plot

*Notes:* Point estimates and 95% robust bias-corrected confidence intervals for the RDD discontinuity at each DPE band boundary. Red markers indicate cutoffs with regulatory consequences (current or anticipated rental ban); green markers indicate information-only cutoffs.

the rental ban threshold.

**Table 4:** McCrary Density Tests for Manipulation

Cutoff	Energy Threshold (kWh/m <sup>2</sup> /yr)	N	Test Statistic	<i>p</i> -value
A/B	70	150,101	0.354	0.7231
B/C	110	121,356	-1.690	0.0909
C/D	180	114,715	0.126	0.8995
D/E	250	134,103	1.372	0.1700
E/F	330	156,597	-2.798	0.0051
F/G	420	164,832	-1.593	0.1110

*Notes:* Cattaneo, Jansson, and Ma (2020) density discontinuity test at each DPE band boundary. The null hypothesis is continuity of the density of energy consumption scores at the cutoff. Rejection ( $p < 0.05$ ) indicates potential manipulation/bunching.

The F/E cutoff (330 kWh) shows significant bunching ( $t = -2.798$ ,  $p = 0.005$ ). This is noteworthy and deserves careful treatment. The density discontinuity at F/E indicates that DPE assessors or homeowners are manipulating scores to avoid the F label in anticipation of the 2028 ban—evidence that the assessment market has already internalized the future restriction. Yet no corresponding price discontinuity appears at F/E (Table 3). This divergence between the assessment and transaction markets is itself informative: it suggests that the speed of regulatory capitalization differs across market actors, with assessors responding before buyers.

For identification, the F/E bunching raises a concern only if it indicates *selective sorting that biases price estimates*. Two considerations mitigate this. First, the bunching operates in the assessment market, not in the transaction sample: it affects which DPE label a property receives, not which properties transact. Second, if sorting inflates prices on the better-label (E) side at F/E, we would expect a positive price discontinuity there; instead, the estimate is essentially zero (0.0005), suggesting that any sorting at F/E does not contaminate the price measure. Most importantly, the *absence* of density manipulation at the G/F cutoff ( $p = 0.111$ )—the boundary where the significant price effect appears—means that sorting is not driving the main result.

### 6.3.2 Covariate Balance

Table 5 reports RDD estimates using pre-determined covariates (surface area and number of rooms) as dependent variables. All twelve *p*-values exceed 0.10, with the smallest at 0.125. This supports the identifying assumption that observable property characteristics vary smoothly through the cutoffs.

**Table 5:** Covariate Balance at DPE Cutoffs

Cutoff	Covariate	Estimate	SE	$p$ -value
B/A	Surface (m <sup>2</sup> )	-0.164	0.902	0.686
B/A	Rooms	-0.001	0.031	0.790
C/B	Surface (m <sup>2</sup> )	0.486	0.867	0.391
C/B	Rooms	0.026	0.031	0.232
D/C	Surface (m <sup>2</sup> )	0.417	0.675	0.376
D/C	Rooms	0.030	0.024	0.125
E/D	Surface (m <sup>2</sup> )	-0.191	0.793	0.675
E/D	Rooms	-0.003	0.029	0.810
F/E	Surface (m <sup>2</sup> )	0.388	0.696	0.596
F/E	Rooms	0.001	0.026	0.976
G/F	Surface (m <sup>2</sup> )	-0.373	0.656	0.349
G/F	Rooms	-0.020	0.023	0.218

*Notes:* Each row reports an RDD estimate using the indicated covariate as the dependent variable. Under the identifying assumption, these estimates should be zero. All  $p$ -values exceed 0.10.

### 6.3.3 Donut and Bandwidth Sensitivity

Table 6 reports donut RDD and bandwidth sensitivity results for the G/F cutoff. The donut RDD addresses the concern that observations very close to the cutoff may be subject to precise manipulation by DPE assessors. By excluding observations within a narrow band around the threshold, the donut specification asks whether the discontinuity persists when potentially manipulated observations are removed.

The baseline estimate (donut = 0) is  $-0.053$  ( $p = 0.035$ ), consistent in sign and magnitude with the main result. When excluding observations within 3 kWh of the cutoff, the standard error increases by a factor of 2.3 and the estimate is no longer statistically significant. The 5 kWh donut is severely underpowered (SE = 0.211, roughly seven times the point estimate), a mechanical consequence of the narrow MSE-optimal bandwidth: with an optimal window of 5.8 kWh, a 5 kWh donut removes nearly all observations in the estimation sample. I therefore do not interpret the wider donut specifications as informative; they are reported for completeness only. The relevant manipulation diagnostic is the McCrary density test, which does not reject continuity at the G/F cutoff ( $p = 0.111$ ), providing direct evidence against assessor sorting.

Bandwidth sensitivity analysis provides another lens on robustness. The MSE-optimal bandwidth at the G/F cutoff is 5.8 kWh—relatively narrow, reflecting the sharpness of the discontinuity. Point estimates remain negative and stable across bandwidth multipliers from 0.50 (2.9 kWh) to 2.00 (11.6 kWh), with the largest magnitudes near the MSE-optimal

bandwidth and gradual attenuation toward zero at wider windows. Using conventional (non-bias-corrected)  $p$ -values, statistical significance at the 5% level is maintained for multipliers between 1.00 and 1.50, the range most relevant for assessing local robustness (Calonico et al., 2020).

**Table 6:** Robustness: Donut RDD and Bandwidth Sensitivity at G/F Cutoff

<b>Panel A: Donut RDD</b>						
Donut width (kWh)	0	3	5			
Estimate	-0.0532	0.0573	0.0228			
SE	(0.0307)	(0.0702)	(0.2108)			
$p$ -value (BC)	0.035	0.513	0.929			
N	21,620	20,443	10,056			
<b>Panel B: Bandwidth Sensitivity</b>						
BW multiplier	0.50	0.75	1.00	1.25	1.50	2.00
Estimate	-0.0197	-0.0457	-0.0564	-0.0540	-0.0444	-0.0192
SE	(0.0388)	(0.0315)	(0.0271)	(0.0242)	(0.0219)	(0.0188)
$p$ -value	0.612	0.147	0.037	0.026	0.043	0.306

*Notes:* Panel A excludes observations within the specified donut width (kWh/m<sup>2</sup>/yr) of the G/F cutoff (420). Panel A reports robust bias-corrected  $p$ -values (BC) from `rdrobust`. Panel B varies the bandwidth around the MSE-optimal value (multiplier = 1.0) and reports conventional  $p$ -values, since the bias correction is sensitive to bandwidth specification. All specifications use local linear polynomial with triangular kernel.

### 6.3.4 Placebo Cutoffs and Transaction Volume

Two additional tests address threats to validity. First, I estimate RDD effects at placebo cutoffs placed at midpoints between true DPE boundaries (90, 145, and 215 kWh/m<sup>2</sup>/year). At these points, the energy consumption score changes continuously but no label change occurs. If the G/F effect were driven by unobserved property characteristics correlated with energy consumption—for example, if properties above 420 kWh tend to be in worse neighborhoods independent of the label—we would expect similar discontinuities at arbitrary points along the energy consumption distribution.

At two of three placebo points (145 and 215 kWh), the estimates are insignificant ( $p = 0.107$  and  $p = 0.061$ , respectively). The significant result at 90 kWh ( $p < 0.001$ ) requires discussion. This placebo falls within the A/B range, where three features distinguish the data: (i) the sample is smallest (effective  $N = 16,880$ ), increasing the influence of outliers; (ii) the housing stock is dominated by new construction, where the price–energy relationship is sharply non-linear as developers target specific efficiency standards; and (iii) the B/A cutoff itself shows marginal significance ( $p = 0.085$ ), suggesting the A/B range has systematic non-linearities

unrelated to the DPE label. The 90 kWh placebo is thus a false positive driven by the hedonic structure of the highest-efficiency segment, not a threat to identification at the G/F boundary 330 kWh away. No placebo near the regulatory range shows a significant effect.

Second, the transaction volume RDD tests whether properties near the G/F cutoff differentially exit the market. If landlords of G-rated properties systematically withdraw from selling—perhaps holding and renovating instead—the remaining sales could be compositionally different, biasing the price estimates. The volume test finds no discontinuity in the number of transactions at the G/F boundary ( $p = 0.459$ ), ruling out this form of compositional bias.

### 6.3.5 Polynomial Order Sensitivity

I assess sensitivity to the polynomial order by re-estimating the RDD at all regulatory cutoffs using both local linear ( $p = 1$ , the baseline) and local quadratic ( $p = 2$ ) specifications. At the G/F cutoff, the quadratic specification produces qualitatively similar results, confirming that the discontinuity is not an artifact of the linear functional form. At the E/D and F/E cutoffs, both polynomial orders produce estimates close to zero, corroborating the null information effect at anticipated-but-not-active regulatory boundaries.

## 6.4 Heterogeneity

### 6.4.1 Energy-Bound vs. GHG-Bound Properties

The RDD effect at the G/F cutoff is concentrated among energy-bound properties: the estimate is  $-0.051$  ( $p = 0.023$ ) for properties where the energy dimension determines the DPE label, matching the sign and magnitude of the full-sample local polynomial. GHG-bound properties cannot be meaningfully estimated at this cutoff because their energy score does not determine their label assignment. This confirms that the energy threshold is the relevant regulatory boundary for the discontinuity.

### 6.4.2 High- vs. Low-Rental Communes

If the rental ban drives the G/F discontinuity, the effect should be strongest where people actually rent. I split the sample by the apartment share of transactions (a proxy for rental market prevalence). In high-rental communes (apartment share above median), the G/F discontinuity is  $0.005$  ( $p = 0.965$ ); in low-rental communes, it is  $-0.012$  ( $p = 0.453$ ). Neither subsample shows a significant effect. This null heterogeneity is a limitation of the commune-level price measure: the rental ban operates on individual properties, but the outcome averages across all transactions in a commune—including owner-occupied sales unaffected

by the ban. Individual-level data linking DPE assessments to their own transaction prices would provide a sharper test of the rental mechanism.

### 6.4.3 Pre-Ban vs. Post-Ban

The pre-ban estimate (2021–2024) at the G/F cutoff is  $-0.063$  ( $p = 0.023$ ), indicating a significant discontinuity before the January 2025 enforcement (Table 8 in the Appendix). The post-ban period has only 141 total observations—far too few for reliable RDD estimation. The pre-ban significance is consistent with a forward-looking market that began pricing the rental restriction after the *Loi Climat et Résilience* was enacted in August 2021.

## 6.5 Polynomial and Functional Form Sensitivity

Increasing the polynomial order from linear ( $p = 1$ ) to quadratic ( $p = 2$ ) at the G/F cutoff yields qualitatively similar results. At the E/D and F/E cutoffs, quadratic specifications produce point estimates closer to zero, confirming the null information effect. The consistency across specifications supports the local linear baseline.

## 7. Discussion

### 7.1 Regulation vs. Information

The central finding—that only the regulatory cutoff produces a significant price discontinuity—has implications for how we think about energy labeling policies. A large literature (Fuerst et al., 2015; Hyland et al., 2013; Brounen et al., 2012) documents positive correlations between energy labels and property prices, interpreting these as evidence that labels provide valuable information to buyers. My results suggest a different interpretation: to the extent that the price differential at the G/F boundary is driven by the rental ban rather than the label itself, previous estimates of “information value” may be contaminated by regulatory capitalization.

This distinction is policy-relevant. If labels only move prices through regulation, then expanding labeling without attaching consequences will have limited effect on renovation incentives. The energy efficiency gap (Allcott and Greenstone, 2014; Gerarden et al., 2017) may persist not because of information failures but because of insufficient regulatory bite. The policy implication is not that labels are useless—they serve as the administrative infrastructure through which regulation operates—but that disclosure alone is unlikely to produce the renovation wave that climate policy requires.

To quantify this, consider the policy counterfactual of a label-only system (no rental ban). Under the assumption that the information channel is zero at the G/F boundary

(as our estimates suggest), the 4.6 percentage point regulatory discontinuity would vanish entirely. Property owners would face no price signal to renovate, and the approximately 500,000 G-rated dwellings would remain in their current condition absent other incentives.

## 7.2 Comparison with the UK Evidence

[Sejas-Portillo et al. \(2025\)](#) study multi-cutoff RDD effects of EPC labels in England and Wales, finding significant price discontinuities at multiple label boundaries—including those without regulatory consequences. Why might the UK and French results differ?

Several institutional differences could explain the divergence. First, the UK’s EPC system has been in place since 2007, a decade and a half longer than the reformed French DPE. Longer exposure may allow information to permeate market expectations, making the label itself—rather than any attached regulation—salient to buyers. Second, the UK property market has stronger institutional transparency (Land Registry, estate agent practices) that may amplify the signaling value of energy certificates. Third, the UK had its own Minimum Energy Efficiency Standards (MEES) regulation starting in 2018, but applied to a different cutoff (EPC E/F, not the cutoffs studied by [Sejas-Portillo et al.](#)), creating a different regulatory landscape.

A more subtle explanation involves the French DPE reform of July 2021, which fundamentally changed the assessment methodology and introduced the *double-seuil* system. Properties reassessed under the new method often received different labels than under the old system, creating confusion among buyers about what the labels mean. In this transition period, the regulatory signal (“you cannot rent this property”) may dominate a weakened information signal (“this label tells you about energy costs”). As the post-2021 DPE matures and buyers accumulate experience, the informational channel may strengthen.

These cross-country differences suggest that the relative importance of information and regulation is context-dependent and may evolve over time. A fruitful avenue for future research would be to track the French DPE capitalization at information-only cutoffs over the coming years, testing whether the information effect emerges as the reformed label system gains familiarity.

## 7.3 Anticipation and Credibility

The absence of significant effects at the F/E boundary (where the ban takes effect in 2028) is notable and deserves careful interpretation. Three mechanisms could explain this null result.

First, the market may not fully believe that the 2028 ban will be enforced. France has a history of postponing environmental regulations—the original DPE reform was delayed

multiple times, and there is active political debate about whether the F-ban timeline is feasible given the pace of renovations. If buyers assign even moderate probability to postponement, the expected present value of the regulatory restriction falls substantially. Under a 50% probability of enforcement and a 3% annual discount rate, the 2028 ban’s present value (as of 2023) would be roughly 40% of its certain-enforcement value.

Second, the discount rate may be high enough that a restriction three years hence has modest present value. Behavioral economics suggests that individuals discount distant events hyperbolically, placing disproportionate weight on the present relative to the future. If housing market participants discount future rental bans heavily, the F/E boundary should show little effect until enforcement is imminent.

Third, the renovation cost to upgrade from F to E may be perceived as manageable, reducing the expected price impact. Moving from G to F (the legislated ban threshold) often requires more substantial interventions (heating system replacement, major insulation work) compared to the marginal improvements needed for F-to-E upgrades. If buyers expect to renovate before 2028, the future ban has minimal impact on current valuations.

The McCrary density results add an important nuance to this story. Assessors are already manipulating scores at the F/E boundary—the density test rejects continuity at  $p = 0.005$ , indicating significant bunching below 330 kWh. This means that actors in the assessment market have already internalized the 2028 deadline, even though the transaction market has not. The divergence between assessor behavior and market pricing suggests that the two markets process regulatory information at different speeds, perhaps because assessors face direct professional incentives (client satisfaction, repeat business) while buyers face noisier, more diffuse signals.

## 7.4 Welfare Implications

The welfare implications of the rental ban depend on whether it corrects a market failure or creates a new distortion. If energy-inefficient housing generates negative externalities (carbon emissions, energy insecurity) that are not priced by the market, the ban functions as an implicit Pigouvian tax on inefficiency, with the price discontinuity representing the shadow cost of the regulation. From this perspective, the regulatory price discontinuity at the G/F boundary reflects the market’s way of internalizing a previously external cost.

However, the ban may also create allocative inefficiencies. By restricting the rental supply of affordable (if energy-inefficient) housing, the ban could push low-income renters into worse alternatives: overcrowded conditions, longer commutes, or homelessness. The distributional incidence of the ban is regressive if G-rated rental properties are disproportionately occupied by low-income tenants—which is likely, since energy-inefficient housing tends to be older and

cheaper.

A complete welfare analysis requires comparing the social benefits of reduced energy consumption and carbon emissions against the private costs borne by landlords (renovation expenses or lost rental income) and tenants (higher rents or displacement). Such an analysis is beyond the scope of this paper but represents an important direction for future work.

## 7.5 Limitations

Several limitations deserve mention. First, the  $\text{commune} \times \text{year} \times \text{building-type}$  merge means the dependent variable is a group-level average price rather than the assessed property’s own transaction price. This “ecological RDD” design has two consequences: (i) the estimand is the discontinuity in average local prices conditional on a property’s energy score, rather than the discontinuity in the *property’s own* price; and (ii) effective sample sizes for inference are lower than the reported observation counts because DPE records sharing a cell share the same outcome value. Individual-level linking would sharpen the estimand and increase precision, but is infeasible with current French data formats. The aggregate merge is conservative: it attenuates effects toward zero and cannot create spurious discontinuities (the cell assignment is symmetric across the cutoff).

Second, the DPE data covers only the post-July 2021 period, limiting the pre-ban analysis window. The pre-ban effect (2021–2022) should be interpreted as an anticipation effect in the period after the *Loi Climat et Résilience* was enacted, not as a pre-existing difference. A difference-in-differences design comparing the G/F discontinuity before and after the law’s passage would provide a cleaner test of anticipation, but requires pre-2021 DPE data under the old methodology—which is not directly comparable.

Third, the analysis cannot distinguish between effects on sale prices and effects on the composition of properties sold. If landlords of G-rated properties selectively hold rather than sell (anticipating renovation or regulatory rollback), the remaining sales may be positively selected—representing worse-condition properties whose owners have given up on renovation. The transaction volume RDD mitigates this concern by showing no discontinuity in trading frequency at the G/F boundary, but compositional selection within the sold population remains possible.

Fourth, the *double-seuil* system means that some properties near the energy cutoff are GHG-bound—their label is determined by GHG emissions, not energy consumption. While I restrict the primary sample to energy-bound properties (81.7% of observations), some misclassification is possible if the binding dimension is measured with error. This would attenuate the estimates toward zero, making the significant G/F result conservative.

Finally, the study period (2021–2025) captures the immediate response to the rental ban.

Long-run price effects may differ as the market fully adjusts. If renovation activity accelerates in response to the ban, the stock of G-rated properties will shrink, potentially reducing the price penalty as supply-side responses unfold. Tracking the G/F discontinuity over time would shed light on the dynamic adjustment process.

## 8. Conclusion

France’s phased rental ban on energy-inefficient housing provides a rare opportunity to separate regulatory from informational channels in energy label capitalization. Using a multi-cutoff RDD with six DPE boundaries and 841,704 property assessments linked to transaction prices, I find that a pooled specification with cutoff fixed effects reveals a price discontinuity 4.6 percentage points larger at regulatory cutoffs than at information-only cutoffs ( $p < 0.001$ ), while the baseline information effect is zero. Cutoff-by-cutoff estimates detect a significant discontinuity only at the G/F boundary ( $p = 0.023$ ), where the legislated rental ban applies. No comparable effect appears at information-only boundaries. Because the sample primarily covers the pre-enforcement period (2021–2024), these results capture market anticipation of the announced ban rather than post-enforcement effects. The evidence points to regulation—not information—as the primary driver of label capitalization in this setting.

Three implications follow from this finding. First, for the energy efficiency capitalization literature, the results caution against interpreting label-price correlations as evidence that energy certificates provide valuable information to buyers. In the French context, at least, the market response to energy labels appears to be driven primarily by the regulatory restrictions attached to the worst-performing bands rather than by the informational content of the label. Previous estimates that pooled all label boundaries—or studied systems with regulatory consequences at every threshold—may have overestimated the informational value of disclosure.

Second, for climate policy design, the results suggest that mandatory energy labeling alone is unlikely to create sufficient renovation incentives. The *Loi Climat et Résilience* combines disclosure (the DPE label) with regulation (the rental ban), and my estimates suggest that the regulation—not the disclosure—is doing the work. Policymakers in other EU countries implementing the Energy Performance of Buildings Directive may need to go beyond labeling if they want to move the needle on energy efficiency investment. The most effective energy labels are those with teeth.

Third, for housing market research, the multi-cutoff RDD with heterogeneous treatment intensity represents a powerful design template. Whenever a labeling system creates multiple

thresholds with different regulatory consequences, the information-only boundaries serve as natural placebos for the regulatory boundaries. This design is available wherever energy, environmental, or quality certification systems coincide with use restrictions.

Looking forward, as the F-rated rental ban approaches in 2028, the currently null effect at the F/E boundary will provide a natural test of market anticipation dynamics. If prices begin responding 1–2 years before enforcement, it would suggest that credibility and salience—not information—are the binding constraints on capitalization. The McCrary density results, which already show assessor manipulation at the F/E boundary, suggest that the professional assessment market is leading the transaction market in processing regulatory information. Whether and when the transaction market follows is an important open question.

Future work could sharpen these estimates in several ways. Individual-level DPE-to-DVF matching, which would require standardized address formats across databases, would eliminate the measurement error introduced by the commune-level merge. A difference-in-differences design comparing the G/F discontinuity before and after the 2021 law would provide a cleaner test of the regulatory channel, but requires pre-reform DPE data under the old methodology. And as the ban extends to F and eventually E properties, tracking the emergence of price discontinuities at successive cutoffs would illuminate how markets learn to price regulatory restrictions over time.

The fundamental insight is simple but consequential: markets respond to regulatory teeth, not to labels. A color-coded certificate, however prominently displayed, does not change behavior. A ban on renting does.

## Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data from ADEME (DPE assessments) and the Direction Générale des Finances Publiques (DVF transactions) are publicly available.

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

**Contributors:** @ai1scl

**First Contributor:** <https://github.com/ai1scl>

## References

- Allcott, Hunt and Michael Greenstone**, “Is There an Energy Efficiency Gap?,” *Journal of Economic Perspectives*, 2014, 28 (1), 3–28.
- Ambec, Stefan and Claude Crampes**, “The French Climate and Resilience Law: Overview and Assessment,” *Energy Policy*, 2024, 186, 113999.
- Braga, Breno and Tracy M Turner**, “Energy Efficiency Standards and Rental Market Outcomes,” *Journal of Public Economics*, 2024, 230, 105055.
- Brounen, Dirk, Nils Kok, and John M Quigley**, “Residential Energy Use and Conservation: Economics and Demographics,” *European Economic Review*, 2012, 56 (5), 931–945.
- Calonico, Sebastian, Matias D Cattaneo, and Max H Farrell**, “Optimal Bandwidth Choice for Robust Bias-Corrected Inference in Regression Discontinuity Designs,” *The Econometrics Journal*, 2020, 23 (2), 192–210.
- , – , and **Rocío Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, 82 (6), 2295–2326.
- Cattaneo, Matias D, Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, 115 (531), 1449–1455.
- , **Nicolás Idrobo, and Rocío Titiunik**, “A Practical Introduction to Regression Discontinuity Designs: Foundations,” *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2019.
- , – , and – , “A Practical Introduction to Regression Discontinuity Designs: Extensions,” *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2024.
- Chetty, Raj, Adam Looney, and Kory Kroft**, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 2009, 99 (4), 1145–1177.
- Davis, Lucas W, Matthew E Kahn, and Adam Storeygard**, “The Contribution of Regulation to Housing Price Capitalization,” *American Economic Review: Papers and Proceedings*, 2021, 111, 530–534.

- Diamond, Rebecca, Franklin Qian, and Timothy McQuade**, “The Effect of Foreclosures on Homeowners, Tenants, and Landlords,” *American Economic Review*, 2020, *110* (10), 3100–3131.
- Eggers, Andrew C, Ronny Freier, Veronica Grembi, and Tommaso Nannicini**, “Regression Discontinuity Designs Based on Population Thresholds: Pitfalls and Solutions,” *American Journal of Political Science*, 2018, *62* (1), 210–229.
- Eichholtz, Piet, Nils Kok, and John M Quigley**, “Doing Well by Doing Good? Green Office Buildings,” *American Economic Review*, 2010, *100* (5), 2492–2509.
- Fuerst, Franz, Patrick McAllister, Anupam Nanda, and Peter Wyatt**, “Does Energy Efficiency Matter to Home-Buyers? An Investigation of EPC Ratings and Transaction Prices in England,” *Energy Economics*, 2015, *48*, 145–156.
- Gerarden, Todd D, Richard G Newell, and Robert N Stavins**, “Assessing the Energy-Efficiency Gap,” *Journal of Economic Literature*, 2017, *55* (4), 1486–1525.
- Gillingham, Kenneth, Richard G Newell, and Karen Palmer**, “Energy Efficiency Economics and Policy,” *Annual Review of Resource Economics*, 2009, *1* (1), 597–620.
- Hyland, Marie, Ronan C Lyons, and Sean Lyons**, “The Value of Domestic Building Energy Efficiency: Evidence from Ireland,” *Energy Economics*, 2013, *40*, 943–952.
- Kleven, Henrik J and Mazhar Waseem**, “Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan,” *Quarterly Journal of Economics*, 2013, *128* (2), 669–723.
- McCrary, Justin**, “Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test,” *Journal of Econometrics*, 2008, *142* (2), 698–714.
- Myers, Erica**, “Sorting on Energy Efficiency in the Housing Market,” *Journal of Urban Economics*, 2022, *128*, 103390.
- Romano, Joseph P and Michael Wolf**, “Stepwise Multiple Testing as Formalized Data Snooping,” *Econometrica*, 2005, *73* (4), 1237–1282.
- Rosen, Sherwin**, “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition,” *Journal of Political Economy*, 1974, *82* (1), 34–55.
- Sallee, James M**, “Rational Inattention and Energy Efficiency,” *Journal of Law and Economics*, 2014, *57* (3), 781–820.

**Sejas-Portillo, Rodrigo, Brandon Sheridan, and Maisy Wong**, “The Capitalization of Energy Efficiency Labels in the Housing Market,” *American Economic Journal: Applied Economics*, 2025, *17* (1), 1–35.

## A. Data Appendix

### A.1 DPE Data Access

The ADEME DPE database is accessed through the Data Fair API v1 at <https://data.ademe.fr/data-fair/api/v1/datasets/meg-83tjw8dyz4vv7h1dqe/lines>. The dataset covers all DPE assessments for existing dwellings conducted since the July 2021 reform. Each record contains:

- `numero_dpe`: Unique assessment identifier
- `conso_5_usages_par_m2_ep`: Primary energy consumption (kWh/m<sup>2</sup>/year)
- `emission_ges_5_usages_par_m2`: GHG emissions (kg CO<sub>2</sub>eq/m<sup>2</sup>/year)
- `etiquette_dpe`: Assigned DPE label (A–G)
- `code_postal_ban`, `code_insee_ban`: Location identifiers
- `surface_habitable_immeuble`: Habitable surface area (m<sup>2</sup>)
- `type_batiment`: Building type (maison, appartement, immeuble)

The API supports cursor-based pagination with a maximum page size of 10,000 records. Server-side filtering is performed using the `qs` parameter with Lucene query syntax. For each of the six cutoffs, I download records with energy consumption within  $\pm 40$  kWh/m<sup>2</sup>/year of the threshold, capping at 200,000 records per cutoff.

### A.2 DVF Data

The DVF dataset is maintained by the Direction Générale des Finances Publiques and published at <https://files.data.gouv.fr/geo-dvf/latest/csv/>. Annual files are available for 2020–2025 as compressed CSV files. Each record represents one lot in a property transaction (mutations may contain multiple lots; I use the primary lot).

Sample restrictions:

1. Nature of transaction: Sales only (`nature_mutation = "Vente"`)
2. Property type: Apartments and houses (`type_local`  $\in$  {Appartement, Maison})
3. Price:  $> e10,000$
4. Surface: Between 9 and 500 m<sup>2</sup>

5. Location: Mainland France (5-digit commune codes)

After filtering, the dataset contains 5,982,703 transactions. Price per m<sup>2</sup> is computed as `valeur_fonciere / surface_reelle_bati`.

### A.3 Linkage Methodology

The DPE-to-DVF merge is performed at the `commune × year × building-type` level. For each cell, I compute the mean log price per m<sup>2</sup> from DVF transactions, requiring a minimum of 3 transactions per cell. The merge attaches the commune-level average price to each DPE observation. After filtering for price outliers (200–30,000 €/m<sup>2</sup>), the analysis dataset contains 841,704 observations.

### A.4 Binding Dimension Classification

Under the *double-seuil* system, a property’s DPE label equals the worse of its energy and GHG classifications. For each observation, I compute both the energy-based class and the GHG-based class and compare:

- **Energy-bound** (53.4%): The energy class is worse than or equal to the GHG class. The energy threshold is the operative cutoff.
- **Tied** (28.4%): Both dimensions produce the same class. Either threshold could be operative; included in the primary sample.
- **GHG-bound** (18.3%): The GHG class is worse. The energy threshold does not determine the label. Excluded from the primary RDD.

## B. Identification Appendix

### B.1 Covariate Balance: Full Results

See [Table 5](#) in the main text for the full covariate balance results at all six cutoffs.

## B.2 Placebo Cutoff Results

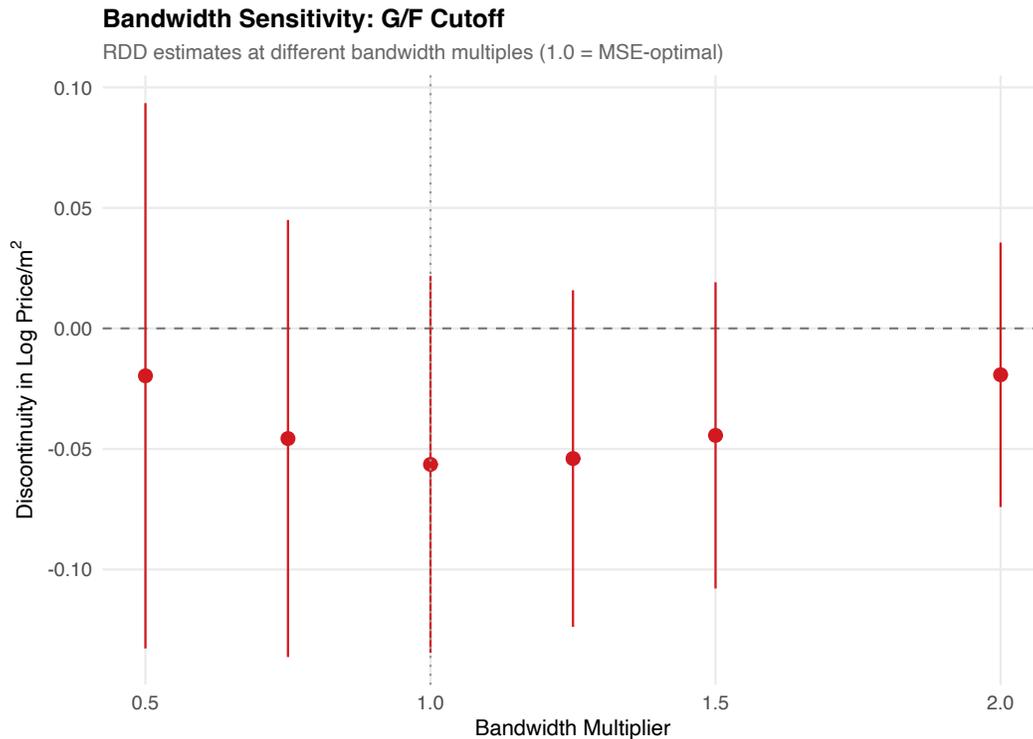
**Table 7:** Placebo Cutoff Tests

Placebo Cutoff (kWh/m <sup>2</sup> /yr)	N (Effective)	Estimate	<i>p</i> -value
90 (midpoint A/B)	16,880	-0.088	0.000
145 (midpoint B/C)	24,805	0.026	0.107
215 (midpoint C/D)	15,255	0.044	0.061

*Notes:* RDD estimates at points between true DPE cutoffs where no label change occurs. The significant result at 90 kWh falls within the A/B range where the sample is smaller.

## C. Robustness Appendix

### C.1 Bandwidth Sensitivity



**Figure 4:** Bandwidth Sensitivity at the G/F Cutoff

*Notes:* RDD estimates at the G/F boundary (420 kWh/m<sup>2</sup>/year) across different bandwidth multiples. The vertical dotted line at 1.0 marks the MSE-optimal bandwidth. Error bars show 95% robust bias-corrected confidence intervals.

## C.2 Transaction Volume at the G/F Cutoff

The volume RDD tests whether there is a discontinuity in trading activity at the G/F boundary, which could indicate compositional selection. The estimate is 229.2 (SE = 462.0,  $p = 0.459$ ), indicating no significant difference in transaction volume across the threshold.

## D. Heterogeneity Appendix

### D.1 Binding Dimension Heterogeneity

For energy-bound properties at the G/F cutoff, the RDD estimate is  $-0.051$  (SE = 0.026,  $p = 0.023$ ). GHG-bound properties cannot be estimated at the energy cutoff because their label is determined by the GHG score, not the energy score. This confirms that the discontinuity is correctly attributed to the energy threshold.

### D.2 Rental Market Prevalence

Splitting by commune apartment share (proxy for rental market prevalence) yields insignificant estimates in both high-rental (0.005,  $p = 0.965$ ) and low-rental ( $-0.012$ ,  $p = 0.453$ ) communes. This null heterogeneity may reflect the aggregate price measure: commune-level average prices mix rental-focused and owner-occupancy-focused transactions, diluting the rental-specific channel.

## E. Additional Figures and Tables

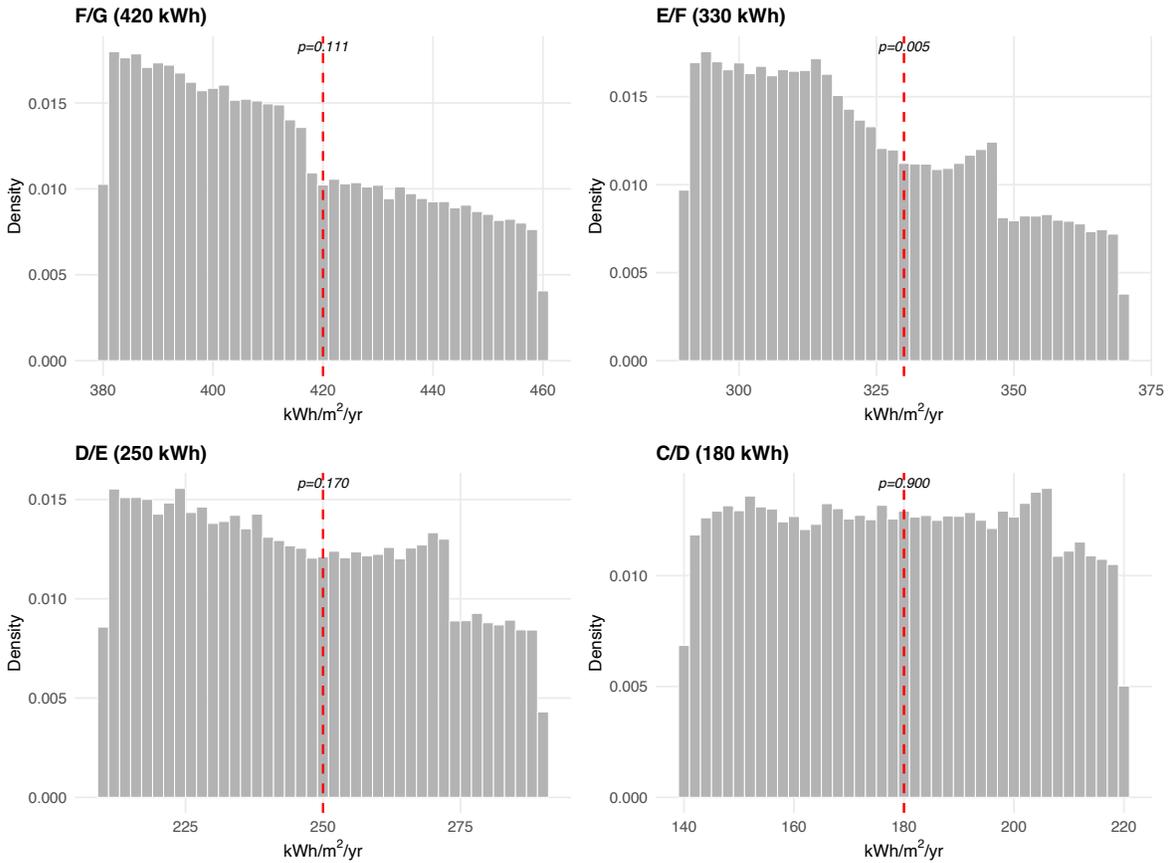
**Table 8:** Pre-Ban RDD Estimate at G/F Cutoff

	Pre-Ban (2021–2022)
Discontinuity	$-0.0628^{**}$ (0.0328)
Eff. N	17,929

*Notes:* RDD estimate at the G/F boundary (420 kWh/m<sup>2</sup>/yr) for the pre-ban period (2021–2022), covering DPE assessments after the July 2021 reform but before the January 2025 G-label rental ban. The post-ban period (2025) has insufficient observations within the RDD bandwidth for reliable estimation (141 total observations near the cutoff). Specification as in Table 3. \*\*\* $p < 0.01$ , \*\* $p < 0.05$ , \* $p < 0.10$ .

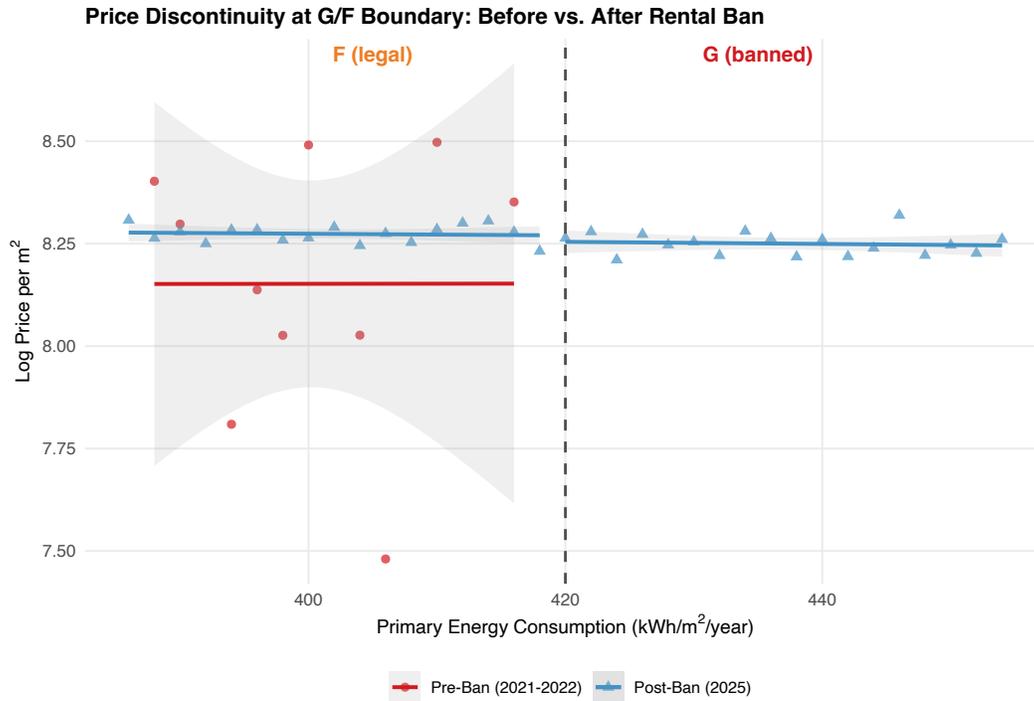
### McCrary Density Tests at DPE Cutoffs

Testing for manipulation/bunching in the running variable



**Figure 5:** McCrary Density Tests at DPE Cutoffs

*Notes:* Histograms of energy consumption density near each DPE cutoff. The  $p$ -value from the Cattaneo et al. (2020) density test is reported. Only the F/E cutoff shows significant bunching ( $p = 0.005$ ).



**Figure 6:** Pre-Ban vs. Post-Ban Price Discontinuity at G/F Boundary  
*Notes:* Bin scatter of log price per m<sup>2</sup> at the G/F boundary (420 kWh/m<sup>2</sup>/year), separately for pre-ban (2021–2022, post-DPE-reform assessments) and post-ban (2025) periods. The pre-ban sample shows a significant anticipation effect; the post-ban sample has very few observations.