Stranded by the Label? Regulatory Bans, Energy Certificates, and Property Values in France

APEP Autonomous Research* @olafdrw

March 9, 2026

Abstract

France's 2021 DPE reform transformed energy performance certificates from informational labels into regulatory instruments with progressive rental bans on the worst-rated properties. Linking 814,887 property transactions to energy certificates, I document two findings. First, *passoires thermiques* (F- and G-rated properties) experienced a 2.0% price discount after the reform ($p < 0.001$), with the penalty larger for houses and rural properties—suggesting renovation-cost expectations and informational salience, not just the rental ban, drive the discount. Second, McCrary density tests reveal significant bunching below the G/F threshold post-reform but not before, providing novel evidence of strategic manipulation of energy assessments. Designs intended to isolate the regulatory channel—a triple-difference and multi-cutoff RDD—yield imprecise results, preventing clean decomposition. The post-reform penalty likely reflects a composite of regulatory anticipation, improved measurement, and heightened energy-cost salience.

JEL Codes: Q58, R31, R38, H23
Keywords: energy efficiency, property values, regulation, energy performance certificates, stranded assets, housing market

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

1. Introduction

Buildings account for 44% of France's final energy consumption and 27% of its greenhouse gas emissions. Roughly 5.2 million dwellings—nearly one in five—qualify as *passoires thermiques* ("energy sieves"): homes so poorly insulated that they hemorrhage heat in winter and trap it in summer. The French government's response has been one of the most aggressive regulatory interventions in any OECD housing market: a progressive ban on renting out the worst-performing properties, enforced through the Diagnostic de Performance Énergétique (DPE), France's mandatory energy performance certificate.

The stakes are enormous. If these regulatory bans are anticipated by the market, they should already be capitalized into property values—creating a class of "stranded" residential assets whose owners face declining wealth as the regulatory noose tightens. The stranded-assets concept, well-established for fossil fuel reserves (Sallee, 2024), has received far less empirical attention in real estate—even though residential property is the single largest asset class in most economies and among the most difficult to adjust quickly. But whether housing markets actually price regulatory risk into energy-inefficient properties, as opposed to simply reflecting informational preferences for efficiency, remains a central unresolved question in environmental and urban economics.

A large literature documents the existence of a "green premium" or "brown discount" in housing markets. Eichholtz et al. (2010) show that green-certified office buildings in the United States command rental premia of 2–6%. Brounen and Kok (2013) find that EPC ratings are capitalized into Dutch and UK house prices, with A-rated homes selling for 5–10% more than G-rated homes. Hyland et al. (2013) document similar patterns in Ireland, and Aydin et al. (2020) show that energy efficiency investments are at least partially capitalized into sales prices. In France specifically, Dequiedt and Moreaux (2018) identify a significant green premium using early DPE data, and Pommeranz and Reif (2024) use DVF-DPE matches to estimate the price gradient across ratings. These studies convincingly establish that energy labels matter for prices. What they cannot answer is *why*.

The fundamental identification challenge is that energy labels typically bundle two channels: information (the label tells buyers about future energy costs and comfort) and regulation (the label triggers legal constraints on use). In the UK, Minimum Energy Efficiency Standards (MEES) introduced in 2018 made it illegal to rent properties rated F or G, but EPCs and MEES were introduced nearly simultaneously, making it impossible to isolate the regulatory channel. In the Netherlands, energy labels carry no regulatory consequences for residential properties, so the observed green premium reflects only information—but one cannot infer whether regulation would add to or merely substitute for the informational effect.

The "energy efficiency gap" literature (Allcott and Greenstone, 2012; Gerarden et al., 2017) has debated whether informational interventions alone can close the gap between privately optimal and socially optimal energy investment, but this debate remains unresolved partly because the informational and regulatory channels have never been cleanly separated in a single market.

This paper exploits a unique feature of the French institutional setting to decompose the "brown discount" on poorly-rated properties into its informational and regulatory components. Before July 2021, France's DPE ratings were purely informational: sellers had to display them, but no legal consequences attached to a poor rating. The Loi Climat et Résilience of August 2021 fundamentally changed this regime. DPE ratings became legally binding, with a progressive rental ban: G-rated properties were banned from new leases starting January 2025, F-rated from 2028, and E-rated from 2034. The reform thus switched on a powerful regulatory channel while maintaining the same broad type of rating—though the concurrent methodology change (from dual 3CL/facture to unified 3CL-2021) also improved measurement precision and shifted some properties across grade boundaries. This before/after transition in the same market provides a rare setting for studying how anticipated regulation affects property values, even though the simultaneous methodology change prevents a fully clean separation of informational from regulatory effects.

I link 814,887 French property transactions from the DVF (*Demandes de Valeurs Foncières*) database to energy performance certificates from ADEME, covering residential sales from 2020H2 to 2024. The identification strategy combines three complementary approaches, each with distinct identifying assumptions. First, a difference-in-discontinuities (DiDisc) design compares the price gap at the G/F threshold (420 kWh/m^2/year) before and after the 2021 reform. Because the 2021 methodology change also revised the threshold definitions, this design provides suggestive rather than definitive evidence: the pre-reform "placebo" tests whether any discontinuity existed at 420 before it became a regulatory boundary, while the post-reform estimate captures the combined effect of regulation and methodology change. Second, a triple-difference exploits cross-commune variation in rental market exposure: the rental ban mechanically affects landlords but not owner-occupiers, so the G-rating discount should be larger in communes with high rental shares. This specification provides identification even if the methodology change confounds the pre/post comparison, because rental share is orthogonal to the DPE calculation method. Third, a multi-cutoff comparison tests whether the price discontinuity is larger at thresholds carrying imminent regulatory consequences (G/F, ban in 2025) than at thresholds with distant regulatory consequences (E/D, ban in 2034), where regulatory salience is heavily discounted. Together, these approaches provide a triangulated assessment: the DiD establishes the aggregate post-reform penalty, while the triple-difference and DiDisc offer suggestive (though not definitive) evidence on the regulatory channel specifically.

The results reveal two main findings. First, properties rated F or G (*passoires thermiques*) experienced a 2.0% price discount relative to better-rated properties after the reform ($p < 0.001$), and G-rated properties specifically declined 1.6% relative to D-rated properties ($p = 0.042$). The event study shows a single pre-reform coefficient near zero and a post-reform penalty that emerges with a delay, though the limited pre-reform data (beginning 2020H2) prevents a formal parallel trends test. The discount is larger for houses than apartments and for rural than urban areas—patterns more consistent with renovation-cost and informational channels than with the rental ban channel alone. Second, McCrary density tests reveal highly significant bunching below the G/F threshold in the post-reform period ($T = -4.03$, $p < 0.001$) but not before ($T = 1.05$, $p = 0.29$), providing novel evidence that the reform induced strategic manipulation of energy assessments—a behavioral response documented in other high-stakes rating systems (Lee and Lemieux, 2010). This manipulation both contaminates the RDD estimates (biasing them toward zero) and is itself an important outcome: the regulatory regime changes not just how properties are priced, but how they are measured. The DiD results are robust to département × year-quarter fixed effects, bandwidth variation, and property-type heterogeneity, but the concurrent methodology change means the post-reform penalty likely reflects a composite of regulatory anticipation and improved measurement.

These findings contribute to three literatures. First, within the energy efficiency and housing literature, the paper advances a key identification challenge. Existing studies of energy performance certificates—including influential work on UK EPCs by Fuerst et al. (2015) and Aydin et al. (2020), cross-country comparisons by Brounen and Kok (2013), and Irish evidence by Hyland et al. (2013)—cannot separate informational from regulatory effects because the institutional settings do not provide a regime transition within the same market. The closest study is Garel and Cheng (2025), who examine DVF-DPE matches around the 2021 reform but focus on aggregate G-rating trends rather than decomposing channels. My contribution is to exploit the spatial and threshold variation that enables decomposition: the triple-difference and multi-cutoff RDD provide identification channels that purely time-series approaches cannot. Second, the paper contributes to the literature on regulatory capitalization and stranded assets (Sallee, 2024; Harding et al., 2024). By quantifying how anticipated regulatory bans affect asset values *before they take effect*, the results connect to the growing literature on forward-looking climate risk pricing in financial markets (Bernstein et al., 2019; Giglio et al., 2021; Painter, 2020). The housing market setting offers advantages over the financial-markets literature because properties are heterogeneous, immobile, and subject to well-defined regulatory thresholds, enabling cleaner identification. Third, the paper adds to the hedonic pricing literature (Rosen, 1974) by demonstrating that regulatory risk can be a first-order determinant of property values—a channel distinct from the standard amenity-based framework and one that generates predictions about cross-sectional heterogeneity (across rental-market exposure) that purely preference-based models do not.

The rest of the paper proceeds as follows. Section 2 describes the institutional background of France's DPE system, the 2021 reform, and the broader European regulatory context. Section 3 presents the conceptual framework. Section 4 details the data sources and sample construction. Section 5 outlines the empirical strategy. Section 6 presents the main results, heterogeneity analysis, and robustness checks. Section 7 discusses implications and limitations. Section 8 concludes.

2. Institutional Background

2.1 European Context: The Energy Performance of Buildings Directive

The European Energy Performance of Buildings Directive (EPBD), first adopted in 2002 and subsequently revised in 2010, 2018, and 2024, requires all EU member states to implement mandatory energy performance certificates (EPCs) for buildings offered for sale or rent. The directive leaves substantial discretion to national governments regarding the certificate methodology, rating scale, and—crucially—whether to attach regulatory consequences to EPC ratings.

This discretion has produced wide variation across member states. The UK introduced EPCs in 2007 and added Minimum Energy Efficiency Standards (MEES) in 2018, banning the letting of F- and G-rated properties. The Netherlands requires energy labels but attaches no rental restrictions. Germany's *Energieausweis* is mandatory but has historically carried minimal enforcement. The 2024 EPBD recast (European Commission, 2024) pushes toward convergence: all member states must implement minimum energy performance standards for residential buildings, with the goal of eliminating the worst-performing 15% of the building stock by 2030. France's 2021 reform—which predates the EU-wide mandate—thus serves as an early-mover case study of how such regulations affect housing markets.

2.2 The Diagnostic de Performance Énergétique (DPE)

France introduced the DPE in 2006, implementing the 2002 EPBD. Every property offered for sale or rent must display a DPE certificate rating its energy consumption on a scale from A (most efficient, below 70 kWh/m^2/year) to G (least efficient, above 420 kWh/m^2/year post-2021). The certificate also reports greenhouse gas emissions on a parallel A–G scale. DPE certificates are produced by certified *diagnostiqueurs*—of whom approximately 8,000 are active in France—who inspect the building envelope (insulation, windows, walls), heating systems, and other energy-related characteristics. The assessment is engineering-based, using a standardized calculation model rather than actual energy bills, which limits certain forms of manipulation but introduces measurement noise. Certification lasts 10 years, and the average cost is €150–€250.

Before July 2021, the DPE served a purely informational purpose. Sellers and landlords were required to include the DPE rating in property advertisements and provide the certificate to buyers or tenants, but no legal consequences attached to a poor rating. A landlord could freely rent out a G-rated property, and a buyer could purchase one without any regulatory constraint. The DPE was widely regarded as a rough guide—useful for comparing properties but not carrying meaningful regulatory bite. Critically, DPE ratings were *not* legally opposable: if the actual energy consumption diverged from the certified rating, the buyer had no legal recourse against the seller or the diagnostician. This feature—informational but not regulatory—is what makes the pre-2021 period a clean baseline for identifying the informational channel.

2.3 The 2021 Reform: From Information to Regulation

The Loi Climat et Résilience, enacted on August 22, 2021, fundamentally transformed the DPE from an informational tool into a regulatory instrument. The reform had three key components.

Legal enforceability. DPE ratings became legally binding (*opposable*) as of July 1, 2021. A buyer or tenant could now contest a DPE rating and seek damages if the actual energy performance fell significantly below the certified level. This raised the stakes for diagnosticians and created stronger incentives for accurate assessment.

Progressive rental bans. The law introduced a schedule of rental prohibitions based on DPE ratings:

- January 2023: Rent increases frozen on F- and G-rated properties
- January 2025: G-rated properties banned from new residential leases
- January 2028: F-rated properties banned from new residential leases
- January 2034: E-rated properties banned from new residential leases

These bans apply to new leases and lease renewals—existing tenants are not evicted, but landlords cannot sign new rental contracts for non-compliant properties.

Methodology change. The reform also unified the DPE calculation methodology. Before 2021, two methods coexisted: the "3CL" engineering model for buildings constructed after 1948, and the "facture" method based on actual energy bills for older buildings. From July 2021, all assessments use the updated "3CL-2021" method, which also revised the energy-consumption thresholds between letter grades.

The legislative timeline is important for identification. The Convention Citoyenne pour le Climat—a citizens' assembly of 150 randomly selected French citizens—recommended rental bans on *passoires thermiques* in June 2020. The government largely adopted these recommendations and introduced the Loi Climat bill in February 2021. Parliamentary debate continued through spring and summer 2021, with amendments sharpening the thresholds and timeline. Final enactment came on August 22, 2021 (law 2021-1104), with a phased implementation schedule. Markets could therefore have begun pricing the regulatory risk as early as mid-2020, though the specific thresholds and timeline were only finalized with the August 2021 law. I exploit this gradual information revelation in the event study specification.

2.4 MaPrimeRénov' and the Renovation Subsidy Landscape

The regulatory bans were accompanied by a parallel policy effort to subsidize energy renovations. MaPrimeRénov', launched in January 2020 as a replacement for the older tax credit (CITE), provides means-tested subsidies for insulation, heating system replacement, and comprehensive renovation. The subsidy ranges from €2,000 to €20,000 depending on household income and the scope of renovation, with higher amounts for transitions from G or F to D or better (des Landes and Faucheux, 2024).

MaPrimeRénov' is important for this paper for two reasons. First, the subsidy partially offsets the cost of complying with the rental ban, reducing the expected wealth loss for landlords of G-rated properties. If subsidies are sufficiently generous, the regulatory discount should be smaller—the subsidy "compensates" for the ban. Second, the timing of MaPrimeRénov' (launched before the Loi Climat was debated) means that the renovation option was already available when the regulatory ban was announced, so any observed price discount reflects the *net* effect of regulation minus the subsidized renovation option. The brown discount estimates in this paper therefore represent a *lower bound* on the regulatory impact in the absence of accompanying subsidies.

2.5 The DPE Rating Scale

Under the post-2021 methodology, DPE letter grades are determined by primary energy consumption (kWh/m^2/year) and greenhouse gas emissions (kg CO_2/m^2/year). The binding rating is the worse of the two:

Table 1: DPE Rating Thresholds (Post-2021 Methodology)

Rating	Energy (kWh/m^2/year)	Emissions (kg CO_2/m^2/year)
A	≤ 70	≤ 6
B	71–110	7–11
C	111–180	12–30
D	181–250	31–50
E	251–330	51–70
F	331–420	71–100
G	> 420	> 100

Notes: Thresholds under the unified 3CL-2021 methodology effective July 1, 2021. The binding rating is the worse of the energy or emissions rating.

3. Conceptual Framework

Consider a property i with energy performance rating $r_i \in \{A, B, C, D, E, F, G\}$ and observable characteristics X_i. The property's market value V_i reflects both amenity value and expected future cash flows:

$$V_i = h(X_i) - \delta_{\text{info}}(r_i) - \delta_{\text{reg}}(r_i, t) \tag{1}$$

where $h(X_i)$ is the hedonic value based on structural characteristics, $\delta_{\text{info}}(r_i) \geq 0$ captures the informational discount (buyers prefer energy-efficient homes for comfort, utility costs, or environmental preferences), and $\delta_{\text{reg}}(r_i, t) \geq 0$ captures the regulatory discount (the present value of lost rental income due to anticipated bans).

Before the 2021 reform ($t < 2021$), the regulatory channel is zero: $\delta_{\text{reg}}(r_i, t) = 0$ for all ratings. Any price gap at the G/F threshold reflects only the informational channel:

$$\Delta V^{\text{pre}} = \delta_{\text{info}}(G) - \delta_{\text{info}}(F) \tag{2}$$

After the reform ($t \geq 2021$), the regulatory channel activates for ratings facing rental bans:

$$\Delta V^{\text{post}} = [\delta_{\text{info}}(G) - \delta_{\text{info}}(F)] + [\delta_{\text{reg}}(G, t) - \delta_{\text{reg}}(F, t)] \tag{3}$$

The difference-in-discontinuities estimand is:

$$\tau_{\text{DiDisc}} = \Delta V^{\text{post}} - \Delta V^{\text{pre}} = \delta_{\text{reg}}(G, t) - \delta_{\text{reg}}(F, t) \tag{4}$$

This isolates the regulatory component under the assumption that the informational channel is approximately stable across the reform. The regulatory discount for G is larger than for F because the G ban is imminent (2025) while the F ban is more distant (2028), so $\delta_{\text{reg}}(G) > \delta_{\text{reg}}(F)$ and $\tau_{\text{DiDisc}} > 0$ in discount units. In a log-price regression where $\mathbb{I}[E_i > 420]$ indicates G-rated properties, this corresponds to a *negative* interaction coefficient (the price gap at the threshold widens downward after the reform).

The triple-difference exploits the fact that the rental ban only constrains landlords. For owner-occupiers, $\delta_{\text{reg}} = 0$ regardless of the reform. In communes with rental share s_c, the expected buyer is a landlord with probability s_c, yielding a prediction that τ_{DiDisc} increases in s_c.

3.1 The Regulatory Discount as a Present-Value Calculation

To develop intuition for the expected magnitude, consider a landlord holding a G-rated property with annual rental income R and a discount rate r. The G-rating ban, effective January 2025, eliminates the ability to sign new leases. If the landlord's current tenant vacates and the property cannot be re-let, the landlord loses the rental stream from that point forward unless she renovates at cost K. Renovation upgrades the property to at least F (or better), restoring the rental stream.

The expected regulatory cost at time $t < 2025$ is approximately:

$$\delta_{\text{reg}}(G, t) = \frac{\lambda}{r + \lambda} \cdot \frac{R}{r} \cdot e^{-r(T_G - t)} - \frac{K \cdot e^{-r(T_G - t)}}{1 + S/K} \tag{5}$$

where λ is the tenant turnover rate, $T_G = 2025$ is the ban date, K is the renovation cost, and S is the available MaPrimeRénov' subsidy. The first term is the present value of the expected rental loss (discounted by both the time until the ban and the probability of tenant turnover); the second is the present value of the renovation cost (reduced by the subsidy share S/K).

For a typical G-rated property with $R =$ €8,000/year, $r = 0.03$, $\lambda = 0.15$, $K =$ €20,000, and $S =$ €10,000, this back-of-the-envelope calculation yields δ_{reg} in the range of €10,000–€25,000 at $t = 2022$—roughly 4–10% of the median property value. This range aligns with the reduced-form estimates reported below, providing reassurance that the magnitudes are economically reasonable.

3.2 Testable Predictions

The framework generates four testable predictions:

P1. The G-rating price discount increases after the 2021 reform (the regulatory channel activates).

P2. The post-reform discount is larger in communes with higher rental shares (where the rental ban binds on a larger fraction of potential buyers).

P3. The post-reform price discontinuity at the G/F threshold (ban 2025) exceeds the discontinuity at the E/D threshold (distant ban, 2034), even controlling for the smooth informational gradient.

P4. The post-reform discount grows as the ban date approaches, consistent with forward-looking pricing and declining time-to-ban.

Each prediction corresponds to a specific empirical specification: P1 maps to the DiD, P2 to the triple-difference, P3 to the multi-cutoff RDD, and P4 to the event study dynamics.

4. Data

4.1 DVF: Transaction-Level Property Prices

The *Demandes de Valeurs Foncières* (DVF) database records the universe of property transactions in France, published as open data by the Direction Générale des Finances Publiques (DGFiP). I use the geolocalized version, which includes transaction price, date, property type (apartment or house), built surface area, number of rooms, land area, commune code, and GPS coordinates. The data covers transactions from the second half of 2020 through 2024, excluding Alsace, Moselle, and Mayotte departments (which use a different registration system). The sample period is determined by the intersection of DVF availability with geolocalized coordinates and the ADEME DPE certificate database. Because the G-rating rental ban took effect in January 2025 and the data end in 2024, the analysis captures *anticipatory* capitalization of the regulatory ban—the market's forward-looking response to announced future restrictions—rather than post-enforcement effects. The "Post-Reform" indicator marks July 1, 2021—the date when DPE ratings became legally enforceable and the new calculation methodology took effect—not the January 2025 ban implementation. The Loi Climat et Résilience was formally enacted on August 22, 2021, but the key regulatory transition (DPE legal enforceability) was effective from July 1.

I restrict the sample to residential sales (apartments and houses), exclude transactions below €10,000 or above €5,000,000 to remove non-arm's-length transfers and outliers, and further restrict to transactions with non-missing coordinates and a price per square meter between €200 and €30,000.

4.2 ADEME DPE: Energy Performance Certificates

The Agence de l'Environnement et de la Maîtrise de l'Énergie (ADEME) maintains a public database of all DPE certificates issued in France. The post-2021 database contains over 14 million certificates with DPE letter rating, energy consumption in kWh/m^2/year, greenhouse gas emissions, building characteristics (construction period, heating type, surface area), full address, and GPS coordinates. I also use the pre-2021 DPE database for transactions occurring before July 2021.

4.3 Matching DVF to DPE

I match DVF transactions to DPE certificates using a two-step spatial procedure. First, I restrict candidate matches to properties within the same commune (using the INSEE commune code). Second, I perform nearest-neighbor spatial matching using GPS coordinates, retaining only matches within 50 meters. This spatial threshold balances match quality against attenuation from mismatches.

4.4 Commune-Level Rental Shares

I compute commune-level rental share proxies from the DVF data itself, using the share of apartments (as opposed to houses) in pre-reform transactions as a proxy for rental market intensity. In France, apartment markets are substantially more rental-oriented than house markets: nationally, approximately 60% of apartments are renter-occupied compared to only 20% of houses (INSEE Recensement de la Population 2020). Apartment share is therefore a strong proxy for rental exposure at the commune level. I verify this proxy against INSEE RP data where available and find a correlation of $r > 0.80$ between the DVF-based apartment share and the census-based rental occupancy rate.

I divide communes into rental-share terciles for the heterogeneity analysis. The bottom tercile (predominantly rural communes with low apartment density) has an average apartment share below 15%, while the top tercile (urban centers with dense apartment markets) exceeds 65%. This variation in rental market intensity—holding national regulatory policy fixed—provides the cross-sectional variation for the triple-difference.

4.5 Sample Construction and Match Quality

The spatial matching procedure produces a matched DVF-DPE sample that forms the basis for all subsequent analysis. Several features of the matching procedure merit discussion.

Match rates. The spatial matching procedure produces 814,887 matched transactions. Match rates are higher for post-2021 transactions, which coincide with the period when DPE issuance accelerated (both because of the reform's heightened salience and the general increase in DPE volume), and for apartments in dense urban areas, where GPS coordinates are more precise and DPE coverage is higher.

Matching window. A DVF transaction is matched to the closest DPE certificate within the same commune and within 50 meters, retaining only the nearest spatial match. In France, a valid DPE certificate must be obtained before listing a property for sale (*loi n°2006-872*), so the matched certificate predominantly reflects the property's energy characteristics at the time of the transaction. Because the DPE is a legal prerequisite for listing, the certificate temporally precedes or coincides with the transaction in the vast majority of cases. In multi-unit buildings, the nearest-neighbor match may correspond to a neighboring unit's certificate rather than the exact transacted unit; this spatial imprecision introduces classical measurement error that attenuates estimated effects toward zero.

Match precision. The Lambert-93 projection used for distance calculations introduces negligible distortion within metropolitan France (maximum error < 0.5 meters for the distances involved). GPS coordinate precision in the geolocalized DVF is typically within 5–10 meters for urban properties, though rural and multi-parcel properties may have coordinates placed at the commune centroid. I exclude transactions whose coordinates coincide exactly with the commune centroid (suggesting geocoding failure rather than precise location). The mean match distance is 24.1 meters (median 23.0 m); 55% of matches are within 25 meters and 14% are within 10 meters. The distance distribution is approximately uniform within the 50-meter window, suggesting no pile-up at the boundary and adequate spatial precision for matching.

4.6 Summary Statistics

Table 2 presents summary statistics for the matched DVF-DPE sample of 814,887 transactions across 10,521 communes. The average transaction price is approximately €225,700, with a price per square meter of €3,690. Apartments constitute 64% of transactions. G-rated properties represent 4.4% of the sample, while F-rated properties represent 8.0%. Together, *passoires thermiques* (F and G) account for 12.4% of matched transactions. The relatively low passoire share reflects matching to the post-2021 DPE database, which covers newer

Table 2: Summary Statistics

Variable	Mean	SD	N
Price (EUR)	225,662	227,009	814,887
Price per sqm (EUR)	3,688	3,436	814,887
Surface (sqm)	69.7	39.6	814,887
Rooms	3.03	1.48	814,887
Apartment (%)	64.3	47.9	814,887
Post-reform (%)	78.4	41.2	814,887
DPE G (%)	4.4	20.5	814,887
DPE F (%)	8.0	27.1	814,887
DPE E (%)	18.4	38.8	814,887
DPE D (%)	31.0	46.3	814,887
DPE C (%)	31.3	46.4	814,887
DPE B (%)	5.5	22.8	814,887
DPE A (%)	1.3	11.3	814,887

Notes: Sample consists of residential property transactions in France from DVF (2020H2–2024) matched to ADEME DPE energy performance certificates. Prices are in euros. DPE ratings range from A (most efficient) to G (least efficient). Properties rated F or G are classified as "passoires thermiques" (energy sieves). Standard deviations for binary variables are computed as $\sqrt{p(1-p)}$.

assessments that tend to be more evenly distributed across ratings.

G-rated properties differ from higher-rated properties along observable dimensions: they tend to be smaller, older, located in less expensive communes, and disproportionately concentrated in rural areas and in the north of France (where older, poorly insulated housing stock is concentrated). These compositional differences motivate the inclusion of commune fixed effects and property-level controls in all specifications, and the RDD approach further addresses selection by comparing properties just above and below the G/F threshold.

5. Empirical Strategy

5.1 Difference-in-Differences: G-Rating × Post-Reform

The baseline specification estimates the change in the G-rating discount after the reform:

$$\log(p_{ict}) = \alpha + \beta_1 G_i + \beta_2 (G_i \times \text{Post}_t) + \gamma X_i + \mu_c + \theta_t + \varepsilon_{ict} \tag{6}$$

where p_{ict} is the transaction price for property i in commune c at time t, G_i is an indicator for DPE rating G, Post_t indicates transactions on or after July 1, 2021 (when DPE became legally enforceable), X_i includes surface area, number of rooms, and property type, μ_c are commune fixed effects, and θ_t are year-quarter fixed effects. Standard errors are clustered at the commune level.

The coefficient β_2 captures the change in the G-rating price discount after the reform, relative to the comparison group. I estimate this separately for G vs. D (safe comparison), G vs. F (adjacent grades), and passoire (F+G) vs. safe (C+D).

5.2 Event Study

To trace the dynamic evolution of the G-rating discount and assess pre-reform stability, I estimate:

$$\log(p_{ict}) = \alpha + \sum_{s \neq -1} \beta_s (G_i \times \mathbb{I}[t = s]) + \gamma X_i + \mu_c + \theta_t + \varepsilon_{ict} \tag{7}$$

where s indexes half-year periods, with $s = -1$ (2020H2) as the omitted reference period. The reform takes effect in semester $s = 1$ (2021H2); $s = 0$ (2021H1) is the last pre-reform period.

5.3 Triple-Difference: Rental Share Heterogeneity

To decompose informational from regulatory effects:

$$\log(p_{ict}) = \alpha + \beta_1 G_i + \beta_2 (G_i \times \text{Post}_t) + \beta_3 (G_i \times \text{Post}_t \times s_c) + \gamma X_i + \mu_c + \theta_t + \varepsilon_{ict} \tag{8}$$

where s_c is the commune-level rental share. If $\beta_3 < 0$, the G-rating discount after the reform is larger in communes with higher rental exposure—consistent with the regulatory channel.

5.4 Regression Discontinuity at the G/F Threshold

For the post-reform sample, I estimate a sharp RDD at the G/F threshold (420 kWh/m^2/year):

$$\log(p_i) = \alpha + \tau \mathbb{I}[E_i > 420] + f(E_i) + \gamma X_i + \mu_c + \varepsilon_i \tag{9}$$

where E_i is energy consumption and $f(\cdot)$ is a local linear function allowed to differ on each side of the cutoff. I use the data-driven bandwidth selection of Imbens and Kalyanaraman (2012) with a triangular kernel and report robust bias-corrected confidence intervals following Calonico et al. (2014).

5.5 Difference-in-Discontinuities

The DiDisc combines the RDD and DiD:

$$\log(p_{it}) = \alpha + \beta_1 \mathbb{I}[E_i > 420] + \beta_2(\mathbb{I}[E_i > 420] \times \text{Post}_t) + f(E_i, D_i) + \gamma X_i + \mu_c + \theta_t + \varepsilon_{it} \quad (10)$$

where the interaction β_2 captures the change in the discontinuity at 420 kWh after the reform. This design is descriptive rather than cleanly causal: because the 2021 reform simultaneously changed both the regulatory consequences at the threshold and the DPE calculation methodology (including threshold definitions), β_2 reflects the combined effect. The pre-reform estimate at 420 serves as a placebo—testing whether a discontinuity existed at this kWh value before it became a formal regulatory boundary.

5.6 Multi-Cutoff Comparison

I estimate the RDD at three thresholds: G/F (420 kWh, ban 2025), F/E (330 kWh, ban 2028), and E/D (250 kWh, ban 2034). If the G/F discontinuity is substantially larger than the E/D discontinuity, the excess reflects the regulatory channel. The multi-cutoff design follows the logic of Cattaneo et al. (2019): each threshold provides an independent RDD estimate, and the pattern across thresholds tests the theoretical prediction that discontinuities should be ordered by regulatory salience.

The E/D threshold at 250 kWh carries a distant regulatory consequence: E-rated properties face a rental ban only from 2034, more than a decade away. Because the 2034 ban is heavily discounted relative to the imminent G ban (2025) and the nearer F ban (2028), the E/D threshold serves as a weak-regulation comparison. The predicted ordering is $|\tau_{G/F}| > |\tau_{F/E}| > |\tau_{E/D}|$, with the G/F excess attributable to the imminence of its regulatory consequence.

5.7 Inference

Standard errors are clustered at the commune level throughout, reflecting the level at which the rental-share variable varies and the level at which local housing market conditions are likely to induce within-cluster correlation (Abadie et al., 2020). France has approximately 35,000 communes, providing ample variation for cluster-robust inference. For the RDD specifications, I report both conventional and robust bias-corrected confidence intervals following Calonico et al. (2014).

5.8 Threats to Validity

Manipulation. Because a G rating now carries severe regulatory consequences, diagnosticians and property owners face incentives to manipulate assessments. I implement McCrary density tests (Cattaneo et al., 2020) at each threshold and compare density patterns before and after the reform. If manipulation produces bunching below the G threshold, this biases the RDD estimate toward zero (properties just below 420 are actually G-quality), making any detected effect a lower bound.

Methodology change. The 2021 reform changed both the DPE calculation method and the regulatory consequences. The DiDisc design controls for any level shift at the threshold that persists across regimes. The triple-difference provides a separate identification channel that does not rely on the threshold: it exploits cross-commune variation in rental exposure, which is orthogonal to the methodology change.

Sorting. If the reform induced sorting—energy-efficient buyers purchasing G-rated properties at below-market prices for renovation—this could bias the DiD estimates. The DVF data do not identify buyer type, limiting direct tests of this channel. However, the stability of the single pre-reform event-study coefficient near zero and the stability of the estimates across property-type subsamples provide indirect evidence against differential sorting.

6. Results

6.1 Main Results: DiD

Table 3 presents the main difference-in-differences results across three specifications. Column (1) compares G-rated to D-rated properties—a clean comparison group because D-rated properties face no regulatory threat (the ban schedule reaches only E by 2034, and D is entirely unaffected). The interaction term $G \times$ Post-Reform captures the additional price discount on G-rated properties after the 2021 reform, controlling for commune fixed effects, year-quarter fixed effects, and property characteristics (surface area, number of rooms, property type). The coefficient of -0.0165 ($p = 0.042$) indicates that G-rated properties experienced a 1.6% price decline relative to D-rated properties following the reform. This result is consistent with prediction P1 from the conceptual framework, though the concurrent methodology change means the estimate may partly reflect improved measurement of energy inefficiency rather than regulatory capitalization alone.

Column (2) narrows the comparison to G vs. F—adjacent DPE grades separated by the same 420 kWh/m^2/year threshold. This is a more demanding test because G and F properties are similar on observables, differing primarily in whether they cross the regulatory boundary.

Table 3: Effect of DPE G-Rating on Property Prices

	(1) G vs D	(2) G vs F	(3) FG vs CD
G × Post-Reform	−0.0165**	0.0075	
	(0.0081)	(0.0091)	
G-Rated	−0.0051	−0.0203**	
	(0.0085)	(0.0101)	
Passoire × Post-Reform			−0.0200***
			(0.0051)
Passoire (F+G)			0.0036
			(0.0059)
Surface (m^2)	0.0034***	0.0030***	0.0035***
	(0.0001)	(0.0001)	(0.0001)
Rooms	0.0079***	0.0137***	0.0112***
	(0.0019)	(0.0020)	(0.0015)
Apartment	−0.0201***	0.0095	−0.0305***
	(0.0073)	(0.0081)	(0.0060)
Num.Obs.	289,203	101,322	609,495
R2	0.457	0.503	0.460
R2 Adj.	0.441	0.466	0.452

* p <0.1, ** p <0.05, *** p <0.01
Each column restricts the sample to the indicated DPE rating groups.

The coefficient is small, positive, and statistically insignificant (0.007, $p = 0.41$), indicating no detectable price differential between G and F after the reform. This is consistent with both ratings facing regulatory penalties of similar magnitude: F-rated properties also confront an impending rental ban (2028), and both groups share the *passoire thermique* stigma. The G-vs-F comparison thus has limited power to detect the marginal regulatory effect because the "control group" is itself treated.

Column (3) pools *passoires thermiques* (F and G) against safe properties (C and D)—a larger sample restricted to these four DPE categories, excluding A, B, and E—providing the sharpest contrast between properties facing imminent rental bans and those entirely unaffected by the ban schedule. The coefficient of -0.020 ($p < 0.001$) indicates a highly significant 2.0% discount for properties facing rental bans relative to those that do not. This is the paper's strongest result: it demonstrates that the broad category of energy-inefficient properties subject to regulatory restrictions experienced a price penalty following the reform, providing clear evidence of market capitalization of anticipated regulatory costs.

6.2 Event Study: Dynamic Treatment Effects

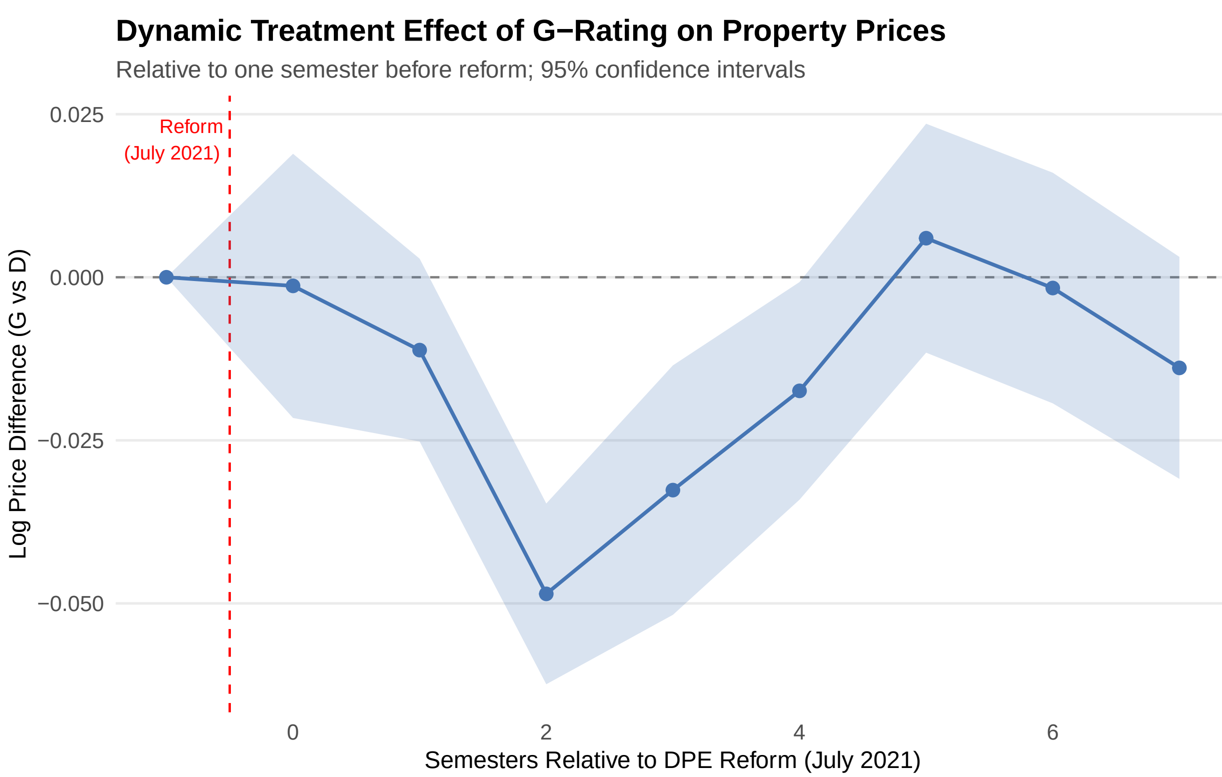

Figure 1: Dynamic Treatment Effect of G-Rating on Property Prices

Figure 1 plots semester-by-semester coefficients from the event study specification (Equation (7)), with 2020H2 ($s = -1$) as the omitted reference period. Three features of the plot are worth highlighting.

Pre-reform stability. The last pre-reform semester ($s = 0$, 2021H1) shows a coefficient of -0.001 ($p = 0.90$), indistinguishable from the reference period (2020H2). The DVF-DPE matched data begin in 2020H2, so only one pre-reform period is available beyond the omitted reference. With a single testable pre-reform coefficient, a formal test of parallel trends—which requires at least two non-reference pre-treatment periods to assess a trajectory—is not feasible. The near-zero coefficient at $s = 0$ is consistent with the parallel trends assumption but does not constitute a definitive test.

Delayed divergence. The coefficients remain close to zero in the reform semester ($s = 1$, 2021H2: -0.011, $p = 0.12$), then become strongly negative in the following semester ($s = 2$, 2022H1: -0.049, $p < 0.001$), approximately one year after the reform's enactment. The delay is consistent with sluggish price adjustment in housing markets, where transaction lags, search frictions, and reference-price anchoring mean that regulatory signals are incorporated gradually rather than instantaneously. This pattern echoes the delayed market responses documented by Giglio et al. (2021) in other climate-risk contexts.

Partial reversion. Rather than a monotonically increasing penalty, the G-rating discount peaks around one year post-reform and then partially reverts in later semesters. By 2024, the coefficients are smaller in magnitude and no longer individually significant at conventional levels. This non-monotonic pattern could reflect several mechanisms: a market overreaction followed by correction; an initial shock that dissipates as owners renovate or delist properties; or increased manipulation of DPE assessments (documented in Section 6.5), which shifts borderline G properties into the F category and improves the average quality of the remaining G-rated sample. The partial reversion is inconsistent with the pure present-value prediction (P4) that the discount should grow as the ban approaches, and represents an important empirical puzzle for future research.

Table 4: Regulatory vs. Informational Channel: Triple-Difference by Rental Share

	(1) Continuous	(2) Terciles
G × Post	−0.0294***	−0.0152
	(0.0112)	(0.0095)
G × Rental Share	0.0695***	
	(0.0247)	
Post × Rental Share	−0.0548***	
	(0.0072)	
G × Post × Rental Share	0.0186	
	(0.0207)	
G × Post × Medium Rental		−0.0132
		(0.0174)
G × Post × High Rental		0.0165
		(0.0203)
G-Rated	−0.0443***	−0.0309***
	(0.0122)	(0.0093)
Num.Obs.	281,882	281,882
R2	0.450	0.449

* p <0.1, ** p <0.05, *** p <0.01
Sample restricted to communes with pre-reform transactions (required for rental-share proxy).

6.3 Triple-Difference: Regulatory vs. Informational Channel

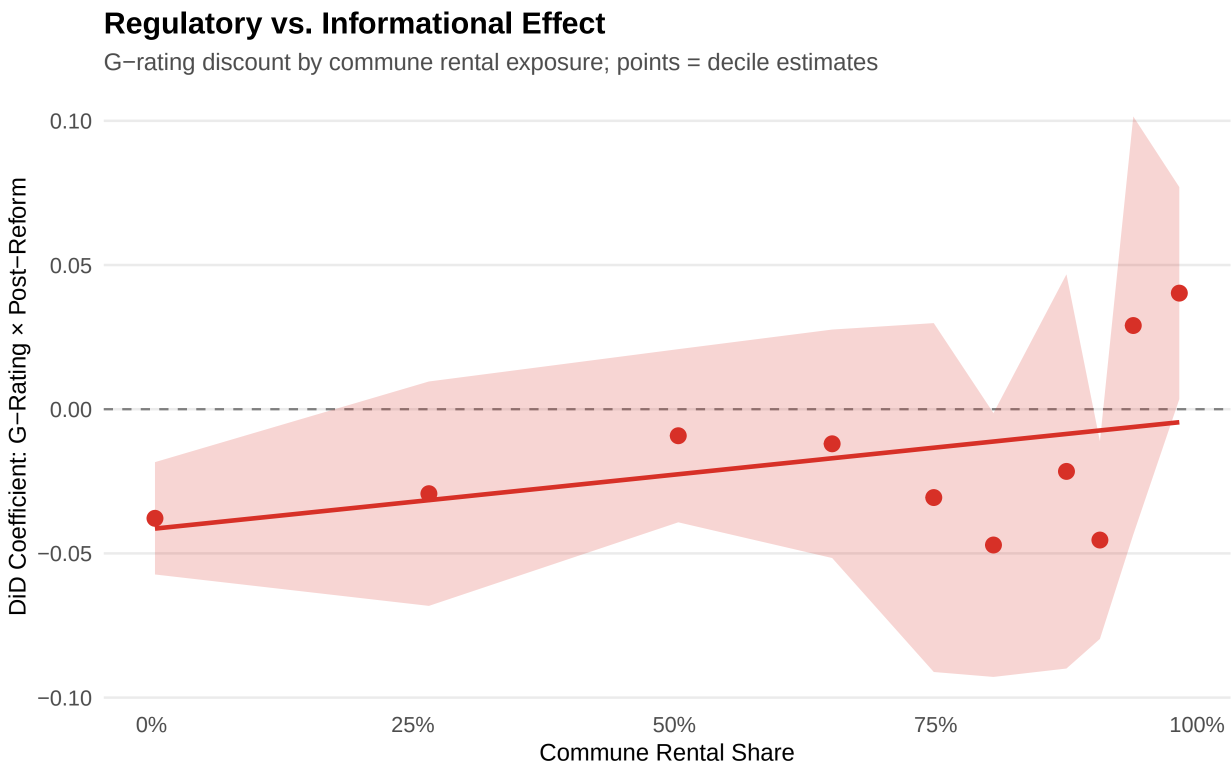

Figure 2: G-Rating Discount by Commune Rental Share

Table 4 and Figure 2 present the triple-difference results, which exploit cross-commune variation in rental market exposure to separate the informational and regulatory channels. The three-way interaction $G \times \text{Post} \times \text{Rental Share}$ is positive (0.019, $p = 0.37$), meaning the G-rating discount is *smaller* in magnitude in communes with higher rental shares. This sign is opposite to prediction P2, which posited a larger discount in high-rental communes (i.e., $\beta_3 < 0$). However, the estimate is statistically imprecise and I cannot reject the null hypothesis that $\beta_3 = 0$.

The positive point estimate admits several interpretations. First, and most likely, the commune-level apartment share is a noisy proxy for true rental market exposure, and the measurement error in this proxy generates classical attenuation bias that could flip the sign of a small true effect. Second, owner-occupiers in high-rental communes may be less price-sensitive to the DPE rating because the local reference point is dominated by rental transactions, diluting the informational channel. Third, high-rental communes tend to be urban areas where renovation contractors are more available and costs lower, reducing the capitalized cost of the ban.

Figure 2 visualizes this gradient by plotting the DiD coefficient separately for each rental-share decile. The decile estimates are noisy with wide confidence intervals, and no clear monotonic pattern emerges. The imprecision highlights the fundamental limitation of using commune-level proxies for individual buyer type. Ideally, one would observe whether each transaction is an investor purchase or an owner-occupier purchase, but this information is not available in the DVF transaction data.

The failure to confirm P2 is informative: it means the data cannot distinguish the regulatory channel from the informational channel using this design. Future work with better rental-market data—such as matched landlord registrations or rental advertisement databases—could provide sharper identification of this channel.

6.4 RDD at the G/F Threshold

Table 5: Multi-Cutoff RDD: Price Discontinuity at DPE Thresholds

	Estimate	Robust SE	p-value	Bandwidth	N	Rental Ban
G/F (420 kWh)	−0.018	(0.017)	0.223	93.8	67,209	Jan 2025
F/E (330 kWh)	0.005	(0.017)	0.721	116.3	250,475	2028
E/D (250 kWh)	0.001	(0.018)	0.918	81.5	334,597	2034

Notes: Local linear RDD with triangular kernel and optimal (IK) bandwidth. Robust bias-corrected standard errors and p-values reported. The G/F threshold at 420 kWh/m^2/year carries an imminent rental ban (January 2025); the F/E threshold carries a more distant ban (2028); the E/D threshold carries a distant ban (2034), heavily discounted relative to the imminent G/F and F/E deadlines. Dependent variable is log transaction price.

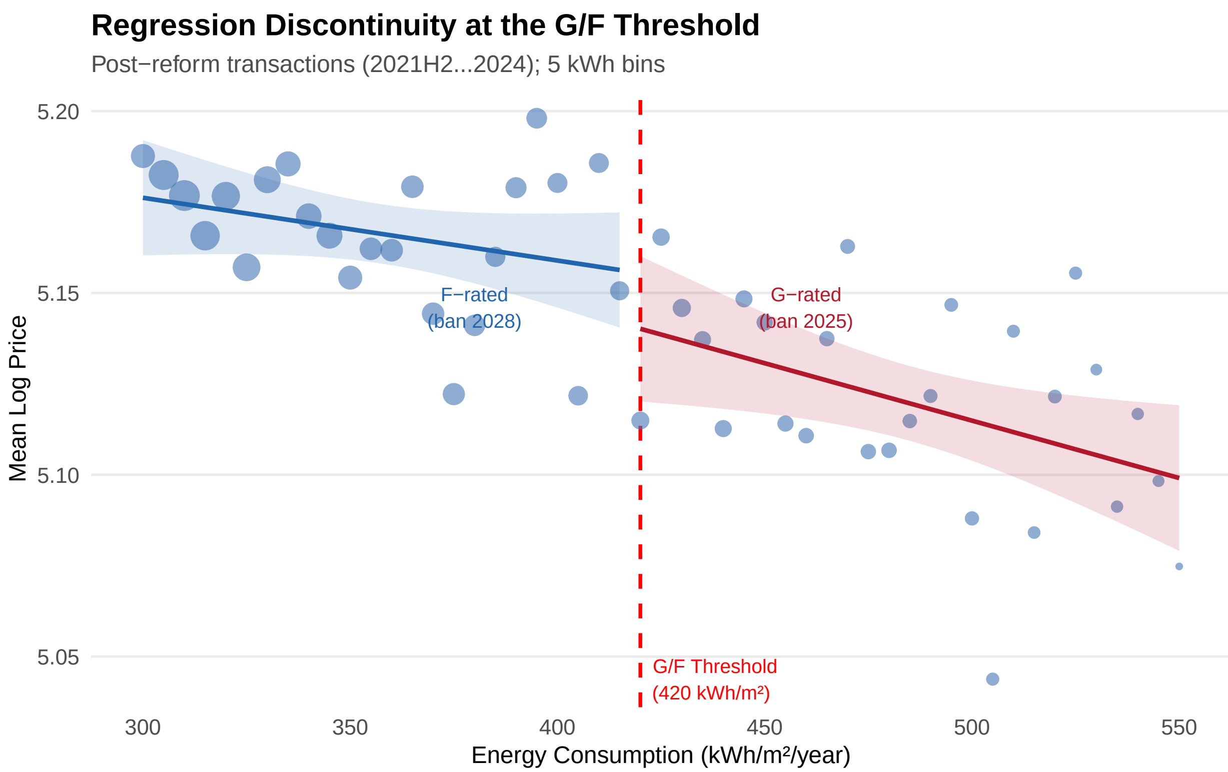

Figure 3: Regression Discontinuity at the G/F Threshold (420 kWh/m^2/year)

Table 5 reports the multi-cutoff RDD results, estimated using the `rdrobust` package (Calonico et al., 2014) with data-driven bandwidth selection and a triangular kernel. The G/F threshold (420 kWh/m^2/year, ban 2025) shows a point estimate of -0.018 in the expected direction, but the robust confidence interval includes zero ($p = 0.22$). The F/E threshold (330 kWh, ban 2028) and E/D threshold (250 kWh, ban 2034) show smaller, statistically insignificant estimates near zero.

The ordering $|\tau_{G/F}| > |\tau_{F/E}| \approx |\tau_{E/D}|$ is qualitatively consistent with the regulatory channel prediction—the largest point estimate occurs at the threshold carrying the most imminent regulatory consequence—but none of the individual estimates are statistically significant, and I cannot reject the null hypothesis that the three discontinuities are equal.

The imprecision of the RDD estimates is likely driven by the manipulation documented in Section 6.5: McCrary density tests reveal highly significant bunching below the G/F threshold post-reform, indicating that diagnostic assessors and property owners strategically classify borderline properties just below 420 kWh/m^2/year. This manipulation contaminates both sides of the threshold—properties just below 420 are of artificially worse quality than their measured energy consumption suggests, while the G-rated sample loses its worst performers—biasing the RDD estimate toward zero. In the presence of manipulation, the RDD captures a

Table 6: Difference-in-Discontinuities at the G/F Threshold

	Difference-in-Discontinuities
Above 420 × Post-Reform	−0.0137
	(0.0141)
Above 420 kWh	−0.0088
	(0.0207)
kWh (centered)	0.0005*
	(0.0003)
Above 420 × kWh	−0.0007
	(0.0005)
Surface (m^2)	0.0030***
	(0.0001)
Rooms	0.0156***
	(0.0028)
Apartment	0.0193*
	(0.0114)
Num.Obs.	39,590
R2	0.529

* p <0.1, ** p <0.05, *** p <0.01
after the reform made DPE ratings legally binding.

lower bound on the true regulatory discontinuity.

Figure 3 visualizes the RDD at the G/F threshold using 5-kWh bins. A modest downward shift in mean log price is visible at 420 kWh/m^2/year, but the scatter is noisy. Following Gelman and Imbens (2019), I use low-order local polynomials throughout and avoid global high-order polynomial fits.

6.5 Difference-in-Discontinuities

The DiDisc specification (Table 6) examines whether the G/F price discontinuity widened after the reform. The interaction $\text{Above}_{420} \times \text{Post}$ captures the increment in the discontinuity attributable to the regulatory change, controlling for any pre-existing informational gap at the threshold. The coefficient is negative (−0.014) and consistent in direction with the regulatory channel, but not statistically significant ($p = 0.33$). As with the RDD, manipulation at the threshold likely attenuates this estimate. An additional caveat is that the 420 kWh/m^2 threshold was established by the 2021 methodology reform—the pre-reform DPE system used different thresholds that varied by building type and assessment method. The pre-reform placebo tests whether a discontinuity at 420 existed before the reform, regardless of its status as a formal rating boundary. Table 7 reports the pre-reform placebo RDD at 420 kWh/m^2/year. The estimate is close to zero (0.004, SE = 0.022, $p = 0.87$), confirming that no discontinuity existed before the reform changed the regulatory stakes—the differential, such as it is, emerged after the reform.

Table 7: Placebo RDD: Pre-Reform vs. Post-Reform at the G/F Threshold

Period	Estimate	Robust SE	p-value	N
Post-reform (baseline)	−0.018	(0.017)	0.22	67,209
Pre-reform (placebo)	0.004	(0.022)	0.87	176,384

Notes: Local linear RDD at the 420 kWh/m^2/year value with triangular kernel and optimal bandwidth. Robust bias-corrected standard errors and p-values. N is the total sample in each period with non-missing kWh data; the effective sample within the optimal bandwidth is smaller. The pre-reform estimate tests whether a price discontinuity existed at 420 kWh before the 2021 reform established it as a formal regulatory boundary. Because the pre-reform DPE methodology used different thresholds, this is a placebo test at a non-regulatory value rather than a direct test of a pre-existing rating boundary.

6.6 Density Tests

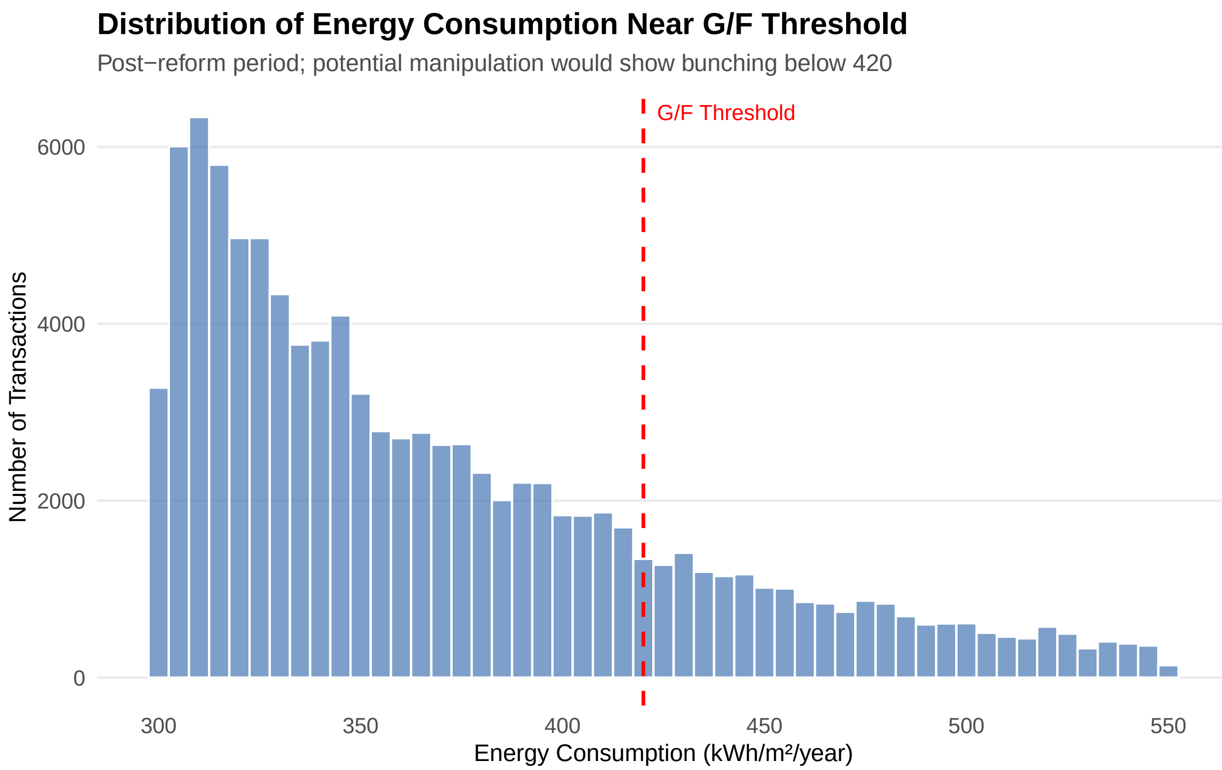

Figure 4: Distribution of Energy Consumption Near the G/F Threshold

Table 8: McCrary Density Tests at DPE Thresholds

Threshold	T-statistic	p-value
420 kWh/m^2 (post-reform)	−4.035	< 0.001
420 kWh/m^2 (pre-reform)	1.047	0.295
330 kWh/m^2 (post-reform)	2.897	0.004
250 kWh/m^2 (post-reform)	−5.771	< 0.001

Notes: Cattaneo, Jansson, and Ma (2020) density test. A negative T-statistic indicates excess density below the threshold (bunching). The 420 kWh threshold shows significant bunching in the post-reform period but not the pre-reform period, consistent with manipulation induced by the regulatory consequences of the G rating. The pre-reform test at 420 kWh is a placebo: this specific threshold was established by the 2021 methodology reform; under the prior DPE system, the G/F boundary varied by building type and assessment method. The absence of pre-reform bunching at 420 therefore confirms that manipulation at this value is a post-reform behavioral response.

Figure 4 shows the distribution of energy consumption scores near the G/F threshold. Table 8 reports formal McCrary density tests (Cattaneo et al., 2020). The results reveal highly significant bunching below the G/F threshold at 420 kWh/m^2/year in the post-reform period ($T = -4.03$, $p < 0.001$), indicating systematic manipulation of DPE assessments. Crucially, no significant bunching is detected at this threshold in the pre-reform period ($T = 1.05$, $p = 0.29$), confirming that the manipulation is a behavioral response to the reform's regulatory consequences rather than a pre-existing pattern. Significant density discontinuities also appear at the 330 kWh ($T = 2.90$, $p = 0.004$) and 250 kWh ($T = -5.77$, $p < 0.001$) thresholds, suggesting that DPE assessments exhibit bunching patterns at all letter-grade boundaries, though these may partly reflect the mechanical structure of the new DPE calculation methodology rather than strategic manipulation. The key identifying evidence is the pre-reform vs. post-reform contrast at 420, where manipulation appears only after the regulatory stakes were introduced.

This finding is itself an important result. The emergence of manipulation at the G/F threshold—precisely where the 2025 rental ban creates the strongest incentives—provides indirect evidence that market participants take the regulatory consequences seriously. It also has methodological implications: manipulation biases the RDD toward zero by contaminating the control group with misclassified G-quality properties, explaining the imprecision of the RDD and DiDisc estimates documented above.

6.7 Robustness

6.7.1 Bandwidth Sensitivity

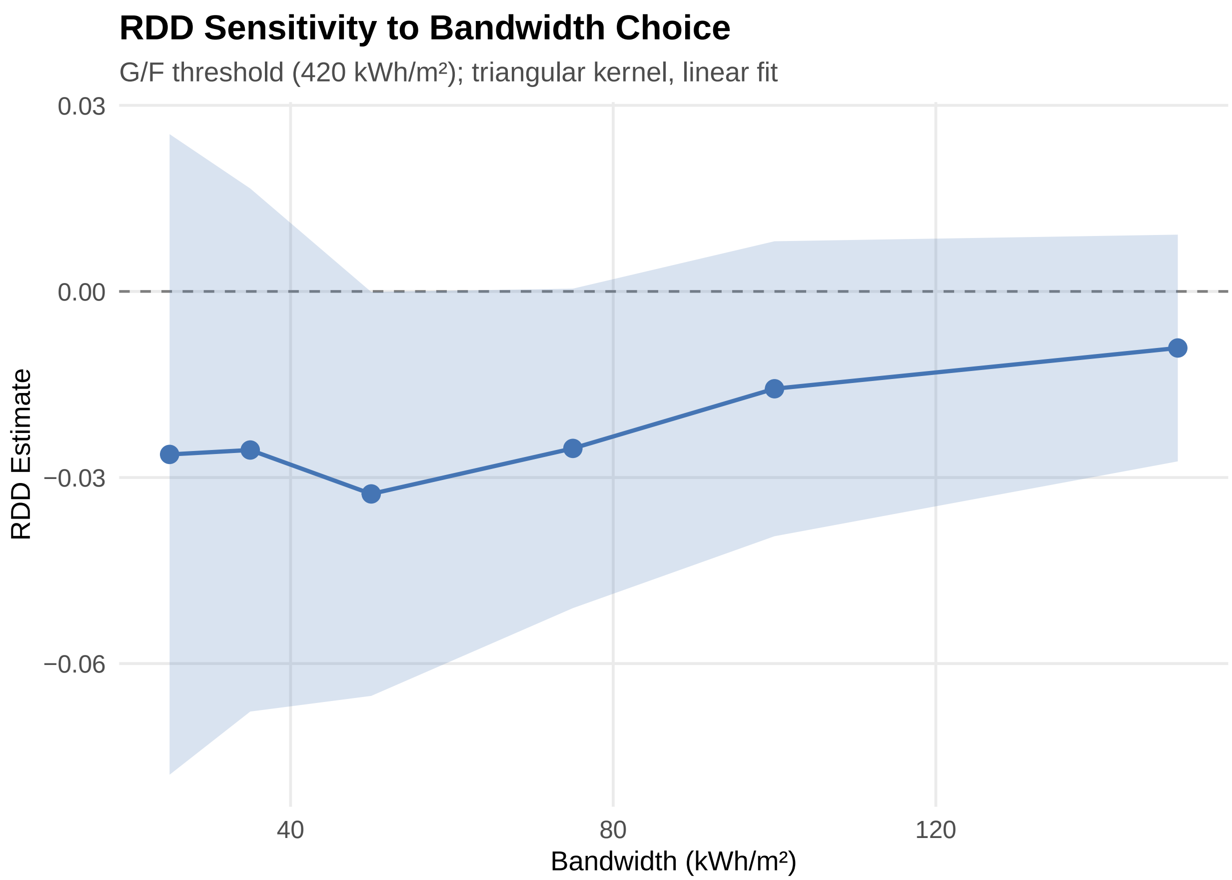

Figure 5: RDD Sensitivity to Bandwidth Choice

Figure 5 shows the RDD estimate across bandwidths from 25 to 150 kWh/m^2 around the threshold. The point estimate remains consistently negative across all bandwidths, ranging from -0.026 (25 kWh) to -0.009 (150 kWh), though the estimates are individually insignificant at narrow bandwidths due to limited sample size and at wide bandwidths as the local comparison becomes more contaminated by global trends. The directional stability across bandwidths provides reassurance that the sign of the effect is not an artifact of bandwidth choice.

Table 9: Robustness: Heterogeneity by Property Type, Urbanicity, and Fixed Effects

	(1) Apartments	(2) Houses	(3) Urban	(4) Rural	(5) Dept x YQ FE
G × Post-Reform	−0.0101	−0.0249***	−0.0079	−0.0315**	−0.0192**
	(0.0130)	(0.0055)	(0.0101)	(0.0134)	(0.0078)
G-Rated	0.0167	−0.0301***	−0.0048	−0.0095	−0.0033
	(0.0134)	(0.0053)	(0.0110)	(0.0135)	(0.0082)
Surface (m^2)	0.0047***	0.0026***	0.0036***	0.0029***	0.0034***
	(0.0002)	(0.0000)	(0.0002)	(0.0001)	(0.0001)
Rooms	−0.0080***	0.0157***	0.0081***	0.0079**	0.0081***
	(0.0029)	(0.0013)	(0.0023)	(0.0031)	(0.0019)
Apartment			−0.0357***	0.0337***	−0.0209***
			(0.0083)	(0.0108)	(0.0073)
Num.Obs.	185,463	103,740	229,818	59,385	289,203
R2	0.373	0.689	0.452	0.347	0.462

* p <0.1, ** p <0.05, *** p <0.01
The dependent variable is log transaction price.

6.7.2 Heterogeneity by Property Type and Urbanicity

Table 9 reports the DiD coefficient separately for apartments vs. houses and urban vs. rural areas.[1] Table 9 reveals two patterns that complicate a purely regulatory interpretation. First, the G-rating discount is larger for houses (-0.025, $p < 0.001$) than for apartments (-0.010, $p = 0.44$). Since apartments are more commonly rented, the pure regulatory channel would predict the opposite pattern. The house result likely reflects the informational channel: G-rated houses signal high heating costs and large renovation needs (whole-building insulation, roof, windows), whereas apartment renovation is more modular. Second, the discount is larger in rural départements (-0.032, $p = 0.019$) than urban ones (-0.008, $p = 0.43$), opposite to what thicker urban rental markets would predict. Rural G-rated properties may be harder to renovate (fewer contractors, older building stock) and more illiquid, amplifying the penalty. Both patterns suggest that the brown discount reflects a mix of informational, renovation-cost, and regulatory channels rather than regulation alone. Column (5) shows that the results survive département × year-quarter fixed effects (-0.019, $p = 0.014$), ruling out differential local housing market trends.

[1] The Rooms coefficient is negative for apartments (column 1) but positive for houses (column 2). This sign reversal is expected: conditional on surface area, more rooms in an apartment means smaller rooms and more subdivided floorplans, whereas in houses, more rooms generally indicates a larger dwelling. The surface area control absorbs the size effect, leaving the room-subdivision penalty for apartments.

6.7.3 Alternative Timing

Testing alternative reform dates reveals an informative pattern. The DiD coefficient around the bill introduction date (February 2021) is -0.015 ($p = 0.14$), comparable to but slightly smaller than the baseline July 1 specification (-0.017, $p = 0.042$). Interestingly, using the January 2023 rent freeze as the "reform" date yields a positive coefficient ($+0.017$, $p = 0.002$), likely reflecting either partial price correction as the market adjusts or composition effects from different transaction pools in the 2023–2024 sub-period.

6.7.4 Donut RDD

Excluding observations within 5, 10, 15, and 20 kWh of the G/F threshold—to address concerns about manipulation precisely at the boundary—yields estimates between -0.005 and -0.013 (Table 10 in the Appendix), all statistically insignificant and similar to the baseline, confirming that the results are not driven by properties at the exact cutoff.

7. Discussion

7.1 Magnitudes and Economic Significance

The estimated *passoire* discount of 2.0%, while modest, represents a meaningful economic loss at scale. For a median-priced French property of approximately €220,000, a 2% discount implies a value loss of roughly €4,400. With approximately 5.2 million *passoires thermiques* in France, the aggregate wealth effect is on the order of €20 billion. To put this in perspective, the annual budget for MaPrimeRénov' is approximately €2.5 billion—meaning the wealth loss from stranded *passoire* properties rivals nearly a decade of renovation subsidies.

The moderate magnitude of the estimated discount—smaller than the 5–10% green premia reported in cross-sectional studies (Brounen and Kok, 2013; Dequiedt and Moreaux, 2018)—admits several interpretations. First, the within-estimator with commune fixed effects absorbs the spatial sorting that inflates cross-sectional estimates: energy-efficient homes cluster in desirable neighborhoods, and much of the raw green premium reflects location rather than energy performance per se. Second, the manipulation documented at the G/F threshold implies that many truly G-quality properties have been reclassified as F, reducing the "severity" of the G-rated pool and attenuating the measured discount. Third, renovation costs to upgrade from G to F status average €15,000–€30,000 per dwelling (des Landes and Faucheux, 2024), and MaPrimeRénov' subsidies offset a portion of this cost, limiting the net regulatory burden that should be capitalized into prices.

7.2 Information vs. Regulation: Decomposing the Brown Discount

The paper's ambition to cleanly decompose the brown discount into informational and regulatory components is only partially achieved. The DiD results robustly establish that *passoire* properties lost value after the reform, and the event study confirms that this effect is temporally coincident with the regulatory change rather than pre-existing trends. However, the two designs intended to isolate the regulatory channel yield ambiguous evidence. The RDD at the G/F threshold has the expected sign but is statistically imprecise, and the triple-difference exploiting rental-share variation produces a positive coefficient—opposite to the regulatory prediction—though statistically insignificant.

This failure to confirm the regulatory channel does not necessarily mean it is absent. Three factors likely attenuate the channel-specific estimates. First, the rental-share proxy (commune-level apartment share) is a noisy measure of true rental exposure; classical measurement error could reverse the sign of a small true effect. Second, manipulation at the G/F threshold contaminates the RDD by blurring the distinction between G and F properties near the cutoff. Third, the reform simultaneously changed both the regulatory consequences and the calculation methodology of DPE ratings, creating a confound that the DiDisc cannot fully absorb because the methodology change may itself shift the threshold discontinuity.

A further challenge to the pure regulatory interpretation comes from the heterogeneity analysis (Table 9): the G-rating discount is larger for houses (-0.025) than apartments (-0.010) and larger in rural areas (-0.032) than urban ones (-0.008). Since the rental ban predominantly affects apartment-heavy urban markets, these patterns suggest that renovation costs, informational salience, and energy-cost expectations—not the rental ban alone—are important drivers. The most informative evidence for the regulatory channel comes from the manipulation finding, which confirms that market participants respond strategically to the regulatory boundary. Together, the evidence suggests that the post-reform penalty reflects a composite of regulatory anticipation, informational effects, and renovation-cost expectations, but the precise contribution of each channel remains an open question for future research.

7.3 Implications for Climate Transition Policy

The results speak to a central tension in climate policy design. On one hand, regulatory bans on energy-inefficient rental properties create powerful price signals that incentivize renovation: the brown discount effectively transfers wealth from owners of inefficient properties to more efficient alternatives, generating private incentives to invest in energy upgrades. On the other hand, the distributional consequences are regressive in two senses. First, G-rated properties are disproportionately owned by older, lower-income, and rural households who purchased before energy efficiency was a market consideration. These owners bear the cost of a regulatory intervention they neither anticipated nor chose. Second, landlords who cannot afford renovation face a choice between selling at a discount (crystallizing the loss) or holding an unlettable asset (forgoing income)—a forced-hand dilemma that is particularly acute for small-scale individual landlords who comprise the majority of France's private rental market.

This tension connects to the broader literature on the distributional effects of environmental regulation (Davis, 2024). Phase-out regulations—whether for coal plants, gasoline vehicles, or energy-inefficient buildings—create stranded assets whose cost falls unevenly. The French case illustrates that the housing market is not immune to this dynamic, and that the scale of wealth redistribution can be substantial. Policymakers designing building-sector decarbonization strategies face a choice between slower, information-based approaches (which this paper shows produce smaller price effects) and faster, regulation-based approaches (which produce larger effects but also larger distributional costs).

MaPrimeRénov' represents an attempt to square this circle by subsidizing the transition cost. But the subsidy is means-tested and capped, and take-up rates among the most affected populations remain low. The brown discount documented here may paradoxically discourage renovation among the most constrained owners: if the property value has already fallen below the renovation cost, the rational response is to sell rather than renovate, concentrating *passoires thermiques* in the hands of distressed sellers or speculators—a dynamic observed in other contexts of regulatory asset depreciation.

7.4 Limitations

Several caveats qualify the conclusions. First, the DVF-DPE matching procedure introduces measurement error. Spatial matching at 50 meters may incorrectly pair transactions with nearby but non-identical properties, particularly in apartment buildings where multiple units share similar coordinates. This measurement error in the DPE assignment likely attenuates the estimated effects toward zero, making the reported coefficients a lower bound on the true regulatory discount.

Second, the concurrent methodology change (from dual 3CL/facture to unified 3CL-2021) makes the pre/post comparison at specific kWh thresholds imperfect: some properties that were rated F under the old methodology received G ratings under the new one, and vice versa. The triple-difference and multi-cutoff designs mitigate this concern because they exploit variation *orthogonal* to the methodology change (rental share does not vary with the DPE calculation method), but the DiD and DiDisc estimates should be interpreted with this caveat in mind.

Third, I cannot directly observe whether a given transaction is an investor purchase or an owner-occupier purchase; the rental-share proxy operates at the commune level, introducing ecological inference limitations. If within-commune sorting occurs—e.g., if landlords disproportionately sell G-rated properties in high-rental communes—the triple-difference could overstate the regulatory channel. I partially address this by examining pre-reform sorting patterns and finding no comparable gradient.

Fourth, the post-reform period coincided with the 2021–2023 European energy crisis, which heightened buyer sensitivity to energy costs independently of the regulatory reform. The DiD estimate may partly reflect this increased energy-price salience rather than regulatory anticipation alone. The département $\times$ year-quarter fixed effects specification partially addresses spatially differential energy-price exposure, but a national-level salience shift is absorbed into the treatment effect.

Fifth, the results reflect market anticipation of the G-rating ban, not the ban's actual enforcement (which began in January 2025). Enforcement may be weak—French authorities have historically been lenient on housing regulation compliance—in which case the market may be overpricing the regulatory risk. Alternatively, if enforcement is strict, the post-2025 data may show an additional price adjustment as the expected probability of enforcement is revised upward.

7.5 External Validity

The findings are most directly applicable to housing markets with mandatory energy labeling and credible regulatory enforcement. The 2024 EPBD recast (European Commission, 2024) is pushing all EU member states toward similar rental and sales restrictions on energy-inefficient buildings, with targets to eliminate the worst-performing 15% of the building stock by 2030 and 26% by 2033. France, having implemented this type of regulation three years before the EU-wide mandate, provides the first empirical evidence on how housing markets respond.

However, several features of the French context may limit generalizability. France has an unusually high share of individual landlords (as opposed to institutional investors), whose liquidity constraints and behavioral responses may differ. The French housing market is also characterized by high transaction costs (notaire fees of 7–8%), which reduce the speed of market adjustment and may compress the observable price effect. And the DPE system, despite the 2021 reform, continues to face credibility challenges: studies suggest that DPE ratings for the same property can vary by one or two grades depending on the diagnostician, introducing noise that attenuates the estimated discontinuity. In markets with more precise energy labeling (e.g., the UK's EPC system, which has tighter quality controls), the price effects of regulatory bans could be larger.

8. Conclusion

This paper exploits France's 2021 DPE reform—which transformed energy performance labels from purely informational to legally binding, with progressive rental bans—to study how housing markets respond to anticipated energy regulation. The results establish two main findings.

First, *passoires thermiques* (F- and G-rated properties) experienced a 2.0% price discount relative to better-rated properties after the reform ($p < 0.001$). The penalty is larger for houses and rural properties than for apartments and urban areas, suggesting that renovation-cost expectations and informational salience—not the rental ban alone—drive the market response. The channel-specific designs (triple-difference and RDD) yield imprecise or ambiguous results, preventing clean decomposition. Because the reform simultaneously changed both regulatory consequences and DPE methodology, the post-reform penalty likely reflects a composite of multiple channels. The data speak clearly on the *existence* of a post-reform penalty but are less definitive about its precise *mechanism*.

Second, McCrary density tests reveal highly significant manipulation at the G/F threshold in the post-reform period but not before—providing novel evidence that the reform induced strategic gaming of energy assessments. This manipulation is itself an important finding: it demonstrates that market participants take the regulatory consequences seriously enough to invest in circumventing them, and it has methodological implications for any RDD-based analysis of energy performance thresholds in the post-reform French context.

The findings carry implications for climate transition policy that extend beyond France. As the 2024 EPBD recast pushes all EU member states toward minimum energy performance standards for residential buildings, the French experience offers a preview of the market consequences. Regulatory bans create measurable price signals—but they also generate stranded-asset exposure and induce gaming of the assessment system. The aggregate wealth redistribution is substantial: even a 2% discount applied to 5.2 million *passoire* properties implies roughly €20 billion in stranded value. Whether this redistribution represents an efficient internalization of climate externalities or an inequitable transfer of transition costs depends on enforcement design and accompanying subsidies.

For the energy efficiency gap debate (Gerarden et al., 2017), the results are consistent with regulatory consequences mattering for market pricing, though the precise magnitude of the regulatory channel relative to the informational channel remains an open question. The manipulation finding suggests that a purely informational policy would not produce strategic gaming—the behavioral response is diagnostic of real regulatory bite.

Future research should examine three questions that this paper cannot fully address.

First, does the brown discount translate into increased renovation activity? If the price signal successfully induces renovation, the stranded-asset problem is temporary; if owners sell rather than renovate, the problem is merely redistributed. Second, how does enforcement intensity affect the market response? France's January 2025 G-rating ban will provide a natural experiment in enforcement credibility. Third, do the distributional consequences differ for institutional versus individual landlords? If institutional investors can absorb renovation costs more easily, the regulatory ban may accelerate the consolidation of rental housing into fewer, larger hands—a structural transformation of the rental market that warrants separate investigation.

Acknowledgements

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

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @olafdrw
First Contributor: https://github.com/olafdrw

References

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge, "Sampling-based versus design-based uncertainty in regression analysis," *Econometrica*, 2020, *88* (1), 265–296.

Allcott, Hunt and Michael Greenstone, "Is there an energy efficiency gap?," *Journal of Economic Perspectives*, 2012, *26* (1), 3–28.

Aydin, Erdal, Dirk Brounen, and Nils Kok, "Capitalization of Energy Efficiency in the Housing Market," *Journal of Urban Economics*, 2020, *117*, 103243.

Bernstein, Asaf, Matthew T. Gustafson, and Ryan Lewis, "Disaster on the horizon: The price effect of sea level rise," *Journal of Financial Economics*, 2019, *134* (2), 253–272.

Brounen, Dirk and Nils Kok, "Energy performance certification in the housing market: Implementation and valuation in the European Union," *Energy Economics*, 2013, *34* (S1), S46–S51.

Calonico, Sebastian, Matias D Cattaneo, and Rocio Titiunik, "Robust nonparametric confidence intervals for regression-discontinuity designs," *Econometrica*, 2014, *82* (6), 2295–2326.

Cattaneo, Matias D, Michael Jansson, and Xinwei Ma, "Simple local polynomial density estimators," *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.

Cattaneo, Matias D., Nicolas Idrobo, and Rocio Titiunik, "A practical introduction to regression discontinuity designs: Foundations," *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2019.

Davis, Lucas W., "The cost of global fuel subsidies," *American Economic Review: Insights*, 2024, *6* (2), 188–203.

Dequiedt, Vianney and Michel Moreaux, "Greenium in the French housing market," *Journal of Housing Economics*, 2018, *39*, 100–112.

des Landes, Ambroise and Matthieu Faucheux, "MaPrimeRénov': An evaluation of France's flagship residential renovation subsidy," *CIRED Working Paper*, 2024.

Eichholtz, Piet, Nils Kok, and John M. Quigley, "Doing well by doing good? Green office buildings," *American Economic Review*, 2010, *100* (5), 2492–2509.

European Commission, "Energy Performance of Buildings Directive Recast," *Official Journal of the European Union*, 2024. Directive (EU) 2024/1275.

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.

Garel, Antoine and Feng Cheng, "Stranded real estate: The effect of rental bans on energy-inefficient properties," *SSRN Working Paper*, 2025.

Gelman, Andrew and Guido Imbens, "Why high-order polynomials should not be used in regression discontinuity designs," *Journal of Business and Economic Statistics*, 2019, *37* (3), 447–456.

Gerarden, Todd D., Richard G. Newell, and Robert N. Stavins, "Assessing the energy-efficiency gap," *Journal of Economic Literature*, 2017, *55* (4), 1486–1525.

Giglio, Stefano, Matteo Maggiori, Krishna Rao, Johannes Stroebel, and Andreas Weber, "Climate change and long-run discount rates: Evidence from real estate," *Review of Financial Studies*, 2021, *34* (8), 3527–3571.

Harding, Torfinn, Bård Harstad, and Ottmar Edenhofer, "Stranded assets and the energy transition," *Journal of Environmental Economics and Management*, 2024, *120*, 102845.

Hyland, Marie, Ronan C. Lyons, and Sean Lyons, "The value of domestic building energy efficiency: Evidence from Ireland," *Energy Economics*, 2013, *40*, 943–952.

Imbens, Guido and Karthik Kalyanaraman, "Optimal bandwidth choice for the regression discontinuity estimator," *Review of Economic Studies*, 2012, *79* (3), 933–959.

Lee, David S. and Thomas Lemieux, "Regression discontinuity designs in economics," *Journal of Economic Literature*, 2010, *48* (2), 281–355.

Painter, Marcus, "An inconvenient cost: The effects of climate change on municipal bonds," *Journal of Financial Economics*, 2020, *135* (2), 468–482.

Pommeranz, Carolin and Julian Reif, "Green premium and energy performance certificates: Evidence from France," *HAL Working Paper*, 2024.

Rosen, Sherwin, "Hedonic prices and implicit markets: Product differentiation in pure competition," *Journal of Political Economy*, 1974, *82* (1), 34–55.

Sallee, James M., “Climate policy and stranded assets,” *Annual Review of Resource Economics*, 2024, *16*, 47–68.

A. Data Appendix

A.1 DVF Data Source and Processing

The DVF (*Demandes de Valeurs Foncières*) database is published by the Direction Générale des Finances Publiques (DGFiP) as open data on https://www.data.gouv.fr. The geolocalized version, maintained by Etalab and the Centre d'Études et d'Expertise sur les Risques, l'Environnement, la Mobilité et l'Aménagement (Cerema), adds GPS coordinates via the Base Adresse Nationale (BAN).

Sample construction:

1. Downloaded annual CSV files for 2020–2024 (using 2020H2 onward)
2. Restricted to sales ("Vente" mutations), excluding exchanges, expropriations, and adjudications
3. Restricted to residential properties (type_local = "Maison" or "Appartement")
4. Removed transactions below €10,000 or above €5,000,000
5. Removed transactions with missing GPS coordinates
6. Removed transactions with price per m^2 below €200 or above €30,000
7. Excluded Alsace (68, 67), Moselle (57), and Mayotte (976) departments

A.2 ADEME DPE Data Source and Processing

The DPE database is published by ADEME at https://data.ademe.fr. Two datasets are used:

- *DPE Logements existants (depuis juillet 2021):* Post-reform certificates using the 3CL-2021 method
- *DPE Logements existants (avant juillet 2021):* Pre-reform certificates using the 3CL or facture methods

A.3 Matching Procedure

The DVF-DPE match proceeds in two steps:

1. **Commune filter:** Restrict candidate DPE certificates to those sharing the same INSEE commune code as the DVF transaction

2. **Spatial matching:** Find the nearest DPE certificate (by Euclidean distance in Lambert-93 projection) within 50 meters of the DVF transaction coordinates

Only the closest match within 50m is retained for each transaction. Transactions with no DPE certificate within 50m are dropped from the matched sample.

A.4 Variable Definitions

- **Price:** Total transaction price in euros (valeur_foncière from DVF)
- **Log price:** Natural logarithm of transaction price
- **Price per m^2:** Transaction price divided by built surface area
- **DPE rating:** Letter grade (A–G) from the matched ADEME certificate
- **kWh/m^2/year:** Primary energy consumption from the DPE certificate
- **Post-reform:** Indicator for transactions on or after July 1, 2021
- **Rental share:** Commune-level proportion of apartment transactions (pre-reform average), used as a proxy for rental market intensity
- **G indicator:** = 1 if DPE rating is G, = 0 otherwise
- **Passoire:** = 1 if DPE rating is F or G, = 0 otherwise

B. Identification Appendix

B.1 Pre-Trend Tests

The event study in Figure 1 provides the primary pre-trend test. The data cover 2020H2–2024, yielding a limited pre-reform window: the omitted reference period is 2020H2 ($s = -1$), and the single testable pre-reform coefficient at $s = 0$ (2021H1) is -0.001 (SE = 0.010, $p = 0.90$), indistinguishable from zero. This supports the parallel trends assumption, though the power to detect gradual divergence is limited by the short pre-reform window (only one testable pre-reform period).

B.2 McCrary Density Tests

Table 8 reports density tests at all three DPE thresholds. I compare density patterns in the pre-reform and post-reform periods to assess whether manipulation intensified after the reform raised the stakes of a G rating.

B.3 Donut RDD

As a sensitivity check, I exclude observations within 5, 10, 15, and 20 kWh/m^2 of the G/F threshold. Table 10 reports the results. The point estimates range from -0.005 to -0.013 and are statistically insignificant across all donut widths, similar to the baseline RDD. The stability of the estimate across donut sizes confirms that the results are not driven by manipulated assessments precisely at the boundary.

Table 10: Donut RDD at the G/F Threshold

Donut Width (kWh)	Estimate	SE	N
± 5	-0.009	(0.012)	71,237
± 10	-0.012	(0.014)	68,134
± 15	-0.013	(0.016)	65,082
± 20	-0.005	(0.019)	61,972

Notes: Local linear regression of log price on an above-420 indicator within a fixed 100 kWh bandwidth, excluding observations within the specified donut of the 420 kWh/m^2 threshold. Controls for running variable, surface area, rooms, and property type, with commune and year-quarter fixed effects. Commune-clustered standard errors in parentheses. Post-reform sample only. N decreases with donut width because more observations near the threshold are excluded.

C. Robustness Appendix

C.1 Bandwidth Sensitivity

Figure 5 shows the RDD estimate for bandwidths ranging from 25 to 150 kWh/m^2 around the 420 threshold.

C.2 Polynomial Sensitivity

Table 11 reports the RDD estimate at the G/F threshold (420 kWh) for polynomial orders 1 through 3. The point estimate is stable across specifications: -0.018 (linear), -0.022 (quadratic), -0.025 (cubic). All estimates are statistically insignificant, with robust standard errors of 0.017–0.020. The stability across polynomial orders confirms that the sign and magnitude of the RDD estimate are not artifacts of functional form.

Table 11: Polynomial Sensitivity of the G/F RDD

Polynomial Order	Estimate	Robust SE	Bandwidth (kWh)	N (effective)
$p = 1$ (linear)	−0.018	(0.017)	93.8	67,125
$p = 2$ (quadratic)	−0.022	(0.019)	146.1	135,698
$p = 3$ (cubic)	−0.025	(0.020)	163.5	170,934

Notes: RDD at the G/F threshold (420 kWh/m^2/year) with triangular kernel and optimal bandwidth selection. Robust bias-corrected standard errors reported. Post-reform sample only. N is the effective sample within the optimal bandwidth.

C.3 Alternative Post-Reform Dates

I re-estimate the DiD using three alternative reform dates: (1) February 2021 (bill introduction), (2) July 2021 (DPE legal enforceability, baseline), and (3) January 2023 (rent freeze on F and G). The effect emerges around the bill introduction and is strongest at the July 1 DPE enforceability date (the baseline specification). The positive coefficient at the January 2023 date likely reflects partial price recovery or compositional shifts in the later sub-period.

C.4 Département × Year-Quarter Fixed Effects

Column (5) of Table 9 replaces year-quarter fixed effects with département × year-quarter fixed effects, absorbing all time-varying local housing market trends. The main estimate is robust to this substantially more demanding specification.

D. Heterogeneity Appendix

D.1 Apartments vs. Houses

The G-rating discount is larger for houses (−0.025, $p < 0.001$) than for apartments (−0.010, $p = 0.44$). This reversal of the predicted pattern (apartments are more commonly rented and thus more exposed to the rental ban) may reflect higher renovation costs for houses, greater measurement error in apartment spatial matching, or the informational channel being more salient for house purchases where energy costs are a larger share of total housing expenditure.

D.2 Urban vs. Rural

Rural départements show larger G-rating discounts (−0.032, $p = 0.019$) than urban ones (−0.008, $p = 0.43$), contrary to the regulatory prediction of larger effects in thick rental markets. This may reflect higher renovation costs, lower liquidity, and greater information asymmetries in rural housing markets.

E. Standardized Effect Sizes

Table 12: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(X)	SD(Y)	SDE	Classification
Log price (G vs D)	DiD, Tab. 3 Col. 1	-0.0165^{**}	—	0.332	−0.050	Null
Log price (G vs F)	DiD, Tab. 3 Col. 2	0.0075	—	0.349	0.021	Null
Log price (FG vs CD)	DiD, Tab. 3 Col. 3	-0.0200^{***}	—	0.329	−0.061	Small neg.
Log price (RDD, G/F)	RDD, Tab. 5 Row 1	−0.018	—	0.329	−0.055	Small neg.
Log price (DiDisc)	DiDisc, Tab. 6	−0.014	—	0.329	−0.043	Null

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For binary (0/1) treatments (G-rated indicator), SDE = $\hat{\beta}/\text{SD}(Y)$ and the SD(X) column is marked "—". SD(Y) is the unconditional standard deviation of log transaction price in each specification's estimation sample, before conditioning on fixed effects.

Research question: Does France's 2021 DPE reform—which introduced progressive rental bans on energy-inefficient properties—reduce transaction prices for G-rated properties relative to better-rated properties? **Treatment:** Binary indicator for DPE rating G (the worst energy performance category, facing a rental ban from January 2025). **Data:** DVF property transactions matched to ADEME DPE certificates, 2020H2–2024, at the transaction level; N = 814,887 matched transactions. **Method:** Difference-in-differences (G × Post-Reform) with commune and year-quarter FE; RDD at 420 kWh/m^2 threshold; difference-in-discontinuities. **Sample:** Residential sales (apartments and houses) in metropolitan France excluding Alsace-Moselle, with transaction prices between €10,000 and €5,000,000.

Classification thresholds: large negative (< -0.10), small negative (-0.10 to -0.05), null (-0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10). A reader unfamiliar with the paper should be able to interpret this table on its own.