Does Rent Control Depress Property Values? Evidence from France's Staggered *Encadrement des Loyers*

APEP Autonomous Research* @olafdrw

March 6, 2026

Abstract

Does limiting rents reduce the sale prices of properties most exposed to rental regulation? I study France's *encadrement des loyers*, which imposed rent ceilings in seven cities between 2019 and 2022, using 451,685 property transactions from *Demandes de Valeurs Foncières* (2020–2024). A triple-difference design compares investment-type properties (studios and small apartments) to owner-occupier properties in five cities with pre-treatment data. The pooled DDD estimate is -0.093 with controls ($p = 0.017$), but the effect is heterogeneous: Bordeaux drives the result (-0.164, $p < 0.001$), while three other cities show null effects. A within-apartment size gradient supports the capitalization mechanism—studios show the largest declines. Randomization inference does not reject the null ($p = 0.46$), reflecting limited power with five treated city groups. The evidence suggests capitalization where regulation binds severely, but cannot establish a general effect.

JEL Codes: R21, R31, R38
Keywords: rent control, property values, capitalization, triple difference, France

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

1. Introduction

When a government caps what landlords can charge, it doesn't just redistribute income from owners to tenants—it revalues the underlying asset. A studio apartment in Bordeaux is worth what its future cash flows are worth, and if those cash flows are capped by law, the studio should sell for less. This capitalization channel is a first-order prediction of standard asset pricing, yet direct evidence on it remains remarkably scarce. France's *encadrement des loyers*—a binding rent ceiling rolled out across seven cities between 2015 and 2022—provides an unusually clean test. The regulation caps rent per square meter, differentiated by neighborhood, building age, room count, and furnishing status. In Paris, where rents had risen 50 percent in real terms over two decades, the stated goal was tenant protection. But the regulation should also depress the sale prices of rental properties relative to owner-occupied ones.

The contribution of this paper is to estimate whether rent control causally depresses property sale prices, exploiting the staggered adoption of the *encadrement des loyers* across seven French cities between 2019 and 2022. I use the universe of residential property transactions recorded in France's *Demandes de Valeurs Foncières* (DVF), an administrative dataset covering every notarized sale in metropolitan France. The identification strategy is a triple-difference (DDD) design: I compare (i) investment-type properties—studios and one- to two-room apartments, which are disproportionately rented—to (ii) owner-occupier-type properties—houses and large apartments, which are primarily owner-occupied—in (iii) cities that adopted rent control versus comparable cities that did not, (iv) before and after adoption. The third difference absorbs city-wide shocks (e.g., a pandemic-driven slump that depresses all property types in treated cities), isolating the differential effect on rental-market-exposed properties.

A critical identification challenge is that the DVF data provide only a five-year rolling window (2020H2–2024), which means Paris (adopted July 2019) and Lille (adopted March 2020) have no pre-treatment observations. I therefore define an "identified sample" of five cities—Plaine Commune, Est Ensemble, Lyon-Villeurbanne, Montpellier, and Bordeaux—that adopted between June 2021 and July 2022, providing genuine pre-treatment data. This identified sample of 451,685 transactions (101,887 treated, 349,798 control) forms the basis for all headline results. Paris and Lille are reported separately as suggestive evidence.

The identified-sample DDD estimate with property controls is -0.093 (SE $= 0.037$, $p = 0.017$): investment-type properties lost approximately 9 percent of their value relative to owner-occupier properties in treated cities after rent control adoption. The effect is similar when measured in price per square meter (-0.103, $p = 0.020$). A within-apartment size gradient reinforces the mechanism: studios and one-room apartments—the units most constrained by rent ceilings—show an additional -0.104 decline ($p = 0.001$) relative to two-room apartments, while three-room apartments show a positive differential ($+0.054$, $p = 0.043$). This monotone gradient is precisely what a capitalization model predicts.

City-level heterogeneity is substantial. Among the five identified cities, Bordeaux drives the result (-0.164, $p < 0.001$), Montpellier is borderline negative (-0.073, $p = 0.062$), and the three remaining cities produce small, insignificant estimates. In supplementary analysis, Paris—though lacking pre-treatment data—shows a striking -0.392 ($p < 0.001$), consistent with severe bindingness in France's tightest rental market. The honest conclusion is that capitalization effects are detectable where rent control is most likely to bind, but the evidence for smaller cities is weak.

This paper contributes to three literatures. First, it advances the empirical literature on rent control, which has focused overwhelmingly on rental market outcomes—tenant mobility (Diamond et al., 2019), housing quality (Early, 2000; Gyourko and Linneman, 1990), and misallocation (Glaeser and Luttmer, 2003; Arnott, 1995b)—rather than on asset prices. The seminal work by Autor et al. (2014) studies spillover effects on property values after the *end* of rent control in Cambridge, Massachusetts, finding that decontrol *raised* nearby property values by 12 percent. Sims (2007) examines a similar decontrol episode. Diamond et al. (2019) documents that San Francisco's rent control reduced rental housing supply by 15 percent as landlords converted to condos. Breidenbach et al. (2022) and Mense et al. (2023) study Berlin's *Mietpreisbremse* and *Mietendeckel* and find negative effects on investment property values. Ahlfeldt and Maennig (2022) estimate that Berlin's rent freeze reduced property values by 8–12 percent. My paper differs from all of these in three respects: it studies the *introduction* rather than the removal of regulation; it exploits staggered multi-city adoption for identification rather than a single policy change; and it uses a triple-difference design that compares property types within cities.

Second, the paper contributes to the literature on capitalization of place-based policies into asset prices (Harding et al., 2012). The insight that regulatory constraints on income streams should be reflected in asset values goes back to Olsen (1972) and Stull (1978). I provide a direct test: if the rental income ceiling binds, properties whose value derives primarily from rental cash flows should depreciate relative to those whose value derives from consumption (owner occupation). The DDD coefficient estimates this differential capitalization.

Third, the paper informs the emerging French literature on the *encadrement des loyers*. Bonnefoy et al. (2021) study early rental market effects in Paris, and Chapelle and Vignolles (2019) examine French housing supply constraints more broadly. Bono and Trannoy (2022) analyze apartment versus house markets. Adda and Pinoli (2023) model misallocation from rent control. To my knowledge, no existing study estimates the causal effect of France's rent control on property sale prices using the DVF administrative data.

2. Institutional Background

2.1 The French Housing Market

France's housing market differs from its Anglo-Saxon counterparts in several institutional features relevant to rent regulation. Approximately 58 percent of French households own their primary residence, compared to 65 percent in the United States and 63 percent in the United Kingdom. The rental sector is divided between private landlords (approximately 23 percent of housing stock) and social housing (17 percent). Private rental housing is concentrated in major cities: in Paris, over 60 percent of housing units are privately rented, compared to less than 20 percent in rural communes.

Property transactions in France are mediated by *notaires* (notaries), public officers with a legal monopoly on real estate transfers. Every sale is notarized, producing a comprehensive administrative record that feeds into the DVF dataset used in this paper. Transaction costs are substantial: buyers typically pay 7–8 percent of the purchase price in *frais de notaire* (notary fees and taxes), creating frictions that reduce speculative turnover compared to markets with lower transaction costs.

The investment property market is characterized by a strong tradition of small-scale individual landlords (*bailleurs particuliers*). Unlike countries where institutional investors dominate rental housing, approximately 95 percent of private rental units in France are owned by individuals, many holding just one or two units. The typical investment property is a studio or one- to two-room apartment in an urban center, purchased with the expectation of steady rental income and capital appreciation. Tax incentives such as the *dispositif Pinel* (2014–2024) and predecessors (Scellier, Duflot) have encouraged apartment purchases in designated zones, further linking the investment property market to rental income expectations.

2.2 Rent Regulation in France

France has a long history of rent regulation, though the modern framework dates to the 1989 *Loi Mermaz-Malandain*, which established a system of reference-based rent adjustments for ongoing tenancies. Under this framework, rents for existing leases can only increase at the rate of the *Indice de Référence des Loyers* (IRL), a consumer price index published quarterly by INSEE. However, for new leases or lease renewals, landlords could historically set rents more freely—a distinction economists call "vacancy decontrol" (Arnott, 1995a).

The key institutional change came in 2014, when the *Loi ALUR* (Accès au Logement et un Urbanisme Rénové) introduced the *encadrement des loyers*—a system of maximum rents per square meter, differentiated by neighborhood (*quartier*), construction vintage (before 1946, 1946–1970, 1970–1990, after 1990), number of rooms, and furnished/unfurnished status. The law authorized local authorities in designated *zones tendues* (tight housing markets) to implement rent ceilings as a function of a median reference rent published annually by the local *observatoire des loyers*.

The mechanism works as follows. Each year, the local *observatoire* publishes a *loyer de référence* (reference rent) per square meter for each combination of neighborhood, vintage, room count, and furnished status. The *loyer de référence majoré* (enhanced reference rent) is set at 120 percent of this median, constituting the legal maximum for new leases. A *loyer de référence minoré* at 80 percent of the median serves as a floor below which tenants can request an increase. Landlords may exceed the *majoré* only by invoking a *complément de loyer* justified by exceptional characteristics (panoramic view, luxury fittings, private garden), but these exceptions are narrowly defined and increasingly challenged in court.

Enforcement relies primarily on tenant initiative. If a tenant believes their rent exceeds the *loyer majoré*, they can challenge the landlord through a *commission départementale de conciliation* (departmental conciliation commission) or in court. Penalties for non-compliance were initially limited to requiring rent reduction to the reference level, but the 2022 reform increased sanctions to include administrative fines of up to €5,000 for individuals and €15,000 for legal entities. Compliance rates have been studied primarily for Paris, where estimates suggest 20–35 percent of new leases exceeded the ceiling in the first years of implementation (Bonnefoy et al., 2021).

2.3 Staggered Adoption

The *encadrement des loyers* was not adopted simultaneously across France. Its rollout followed a staggered pattern driven by legal challenges, political will, and administrative capacity:

- **Paris** (July 1, 2019): After an initial implementation in 2015–2017 that was struck down by the Tribunal administratif, Paris reinstated rent control under a new legal basis (the 2018 *Loi ELAN*).
- **Lille, Hellemmes, and Lomme** (March 1, 2020): The Lille metropolitan area (Métropole Européenne de Lille) adopted rent control for its three central communes.
- **Plaine Commune** (June 1, 2021): An *intercommunalité* of nine communes in the northern Paris suburbs (Aubervilliers, Épinay-sur-Seine, L'Ile-Saint-Denis, La Courneuve, Pierrefitte-sur-Seine, Saint-Denis, Saint-Ouen, Stains, Villetaneuse).

- **Lyon and Villeurbanne** (November 1, 2021): The two central communes of the Lyon metropolitan area.
- **Est Ensemble** (December 1, 2021): Nine communes in the eastern Paris suburbs (Bagnolet, Bobigny, Bondy, Le Pré-Saint-Gervais, Les Lilas, Montreuil, Noisy-le-Sec, Pantin, Romainville).
- **Montpellier** (July 1, 2022): The first city outside the Paris and Lyon metropolitan areas to adopt rent control.
- **Bordeaux** (July 15, 2022): The central commune of the Bordeaux metropolitan area.

This staggered adoption provides the temporal variation that the identification strategy exploits. Critically, the five cities adopting from June 2021 onward have pre-treatment data in the DVF window, while Paris and Lille do not.

2.4 How Rent Control Binds

The *encadrement* sets a *loyer de référence majoré* (enhanced reference rent) equal to 120 percent of the neighborhood median, beyond which new lease rents cannot be set without justification of exceptional property characteristics. Early evidence suggests the regulation binds primarily at the top of the rent distribution—luxury studios and small apartments in central neighborhoods where market rents significantly exceed the reference (Bonnefoy et al., 2021). This is precisely the segment of the market most relevant for investment properties.

For a property investor, the regulation truncates the upper tail of expected future rental income. Standard asset pricing implies that the sale price of a rental property equals the present value of expected net rents. If the regulation reduces expected rents by ΔR per period, the property's value should fall by approximately $\Delta R/r$, where r is the discount rate. Investment-type properties—small apartments likely to be rented—should therefore lose value relative to owner-occupier properties, whose value derives from consumption rather than rental income.

3. Conceptual Framework

Consider two types of residential property in a city c: *investment-type* properties (I), which are purchased primarily to generate rental income, and *owner-occupier-type* properties (O), which are purchased primarily for consumption. The sale price of an investment property reflects the present value of expected future net rental income:

$$P_{ct}^I = \sum_{s=0}^{\infty} \frac{\mathbb{E}_t[R_{c,t+s}^I - M_{c,t+s}^I]}{(1+r)^s} \tag{1}$$

where $R_{c,t+s}^I$ is gross rent, $M_{c,t+s}^I$ is maintenance and management costs, and r is the discount rate. The sale price of an owner-occupier property reflects the present value of imputed consumption services:

$$P_{ct}^O = \sum_{s=0}^{\infty} \frac{\mathbb{E}_t[V_{c,t+s}^O]}{(1+r)^s} \tag{2}$$

where $V_{c,t+s}^O$ is the flow value of housing services to the owner-occupier.

When rent control is introduced in city c at time τ_c, it imposes a ceiling $\bar{R}_c$ on market rents. If $\bar{R}_c < R_{c,\tau_c}^I$ for a subset of investment properties—that is, if the control binds—then $\mathbb{E}_{\tau_c}[R_{c,t+s}^I] < \mathbb{E}_{\tau_c-1}[R_{c,t+s}^I]$ for $s \geq 0$, and:

$$\Delta P_c^I = P_{c,\tau_c}^I - P_{c,\tau_c-1}^I < 0 \tag{3}$$

The key prediction is differential: rent control should not directly affect the consumption value of owner-occupier properties (unless through general equilibrium channels like neighborhood quality changes). Therefore:

$$\Delta P_c^I - \Delta P_c^O < 0 \tag{4}$$

This is the object that the triple-difference design estimates: the differential change in log prices between investment and owner-occupier properties, in treated versus control cities, after rent control adoption.

3.1 Predictions

The model generates several testable predictions about the cross-sectional pattern of price effects.

Prediction 1: Size gradient. The effect should be larger for property types where the regulation is more likely to bind. Studios and one-room apartments in high-rent neighborhoods face the tightest constraints under the *encadrement*, because their rents per square meter are typically highest (landlords charge a premium per square meter for small units) and the reference rents are calculated separately for each room-count category. The prediction is a monotonically decreasing (more negative) DDD coefficient as property size decreases.

Prediction 2: City-level heterogeneity. The effect should be larger in cities where the gap between market rents and reference rents is wider. Paris, with the tightest rental market in France, should show the strongest capitalization effect. Cities where the *encadrement* ceiling is above or near market rents should show no effect.

Prediction 3: Owner-occupier placebo. The effect should be absent or attenuated for owner-occupier-type properties (houses, large apartments) that are rarely rented and hence unaffected by the rent ceiling. In the event study, the post-adoption coefficients for owner-occupier properties should be near zero.

Prediction 4: Volume composition. If rent control reduces the attractiveness of investment properties, we may observe a shift in the composition of transactions: a declining share of investment-type sales in treated cities relative to control cities, as investors hold (avoiding crystallizing losses) or convert to owner-occupation.

3.2 Limitations of the Framework

The stylized model abstracts from several real-world features. First, it treats the regulation as a permanent, credible constraint. In practice, the *encadrement* has faced legal challenges (it was struck down in Paris in 2017, then reinstated in 2019) and compliance is imperfect. If investors assign positive probability to the regulation being weakened or repealed, the capitalization effect would be smaller than the model predicts.

Second, the model treats investment and owner-occupier properties as distinct markets. In reality, the boundary is porous: a studio can be purchased for rental income or for personal use, and market participants arbitrage between the two. If rent control causes investors to exit the studio market and owner-occupiers to enter, the equilibrium price effect may be smaller than the partial-equilibrium prediction.

Third, the model ignores the option value of conversion. A landlord who faces reduced rental income can convert the property to owner-occupation, sell to an owner-occupier, or (in some cases) convert to short-term rental. These outside options provide a floor on the property's value that may limit the extent of capitalization. Diamond et al. (2019) document that San Francisco landlords responded to rent control by converting rental units to condominiums, reducing the rental stock by 15 percent. Similar conversions in Paris would attenuate the measured price decline for investment-type properties.

4. Data

4.1 Demandes de Valeurs Foncières (DVF)

The primary data source is the *Demandes de Valeurs Foncières* (DVF), an administrative dataset recording the universe of notarized property transactions in metropolitan France. Published by the Direction Générale des Finances Publiques via `data.gouv.fr`, DVF covers every residential and commercial property sale, with information on transaction price (*valeur foncière*), property type, surface area, number of rooms, location (commune, postal code), and date.

I use DVF data covering 2020H2 through 2024, the most recent five-year rolling window available. The data are downloaded as pipe-delimited text files from the official `data.gouv.fr` portal, one file per year (plus a partial 2020 file covering the second half). Each file is parsed, standardized, and combined into a single dataset of 5.3 million residential transactions across 92 départements and over 33,000 communes, after excluding Alsace-Moselle (départements 57, 67, and 68), which are governed by a different land registry system (*livre foncier*) and excluded from DVF.

The DVF data have several advantages for this study. First, they are exhaustive: every notarized sale is recorded, eliminating sample selection concerns that plague survey-based studies. Second, the data include both the transaction price and detailed property characteristics—type, surface area, room count—that are essential for the investment/owner-occupier classification. Third, the geographic identifiers (5-digit INSEE commune codes) allow precise treatment assignment at the commune level, matching the administrative units at which the *encadrement des loyers* is implemented.

A key limitation is that DVF provides a five-year rolling window rather than a historical archive. The Direction Générale des Finances Publiques publishes only the most recent five years of transactions, and older data are not publicly accessible. This means I observe no transactions before mid-2020—after Paris (July 2019) and Lille (March 2020) had already adopted rent control. For these two cities, I have only post-treatment data and cannot estimate pre-treatment trends. This constraint is fundamental to the identification strategy and motivates the "identified sample" approach described in Section 5.

A second limitation is that DVF records the property sale price (*valeur foncière*), not the rental income. I cannot directly observe whether a property is rented, at what price, or whether the *encadrement* ceiling binds. The property-type classification serves as a proxy for rental market exposure, but it introduces measurement error. DVF also does not record the buyer's identity or stated purpose (investment vs. owner-occupation), preventing a direct classification.

4.2 Sample Construction

I restrict the sample to residential sales (houses and apartments) with transaction prices between €10,000 and €50,000,000 and building surface area between 5 and 5,000 m^2. Transactions with implausible price-per-square-meter values (below €200/m^2 or above €30,000/m^2) are dropped. The full sample includes 45 treated communes across 7 cities and 20 control cities with populations above 100,000 that never adopted rent control during the sample period (Toulouse, Nantes, Nice, Rennes, Rouen, Toulon, Saint-Étienne, Le Havre, Reims, Dijon, Angers, Clermont-Ferrand, Tours, Limoges, Amiens, Perpignan, Brest, Besançon, Orléans, and Caen). This yields 621,351 transactions: 271,553 in treated communes and 349,798 in control cities.

The **identified sample**—excluding Paris and Lille, which have no pre-treatment data—comprises 451,685 transactions: 101,887 in treated communes across five city groups and 349,798 in control cities. All headline results use this sample.

4.3 Property Type Classification

I classify each transaction as "investment-type" or "owner-occupier-type" based on observable characteristics:

- **Investment-type**: Apartments with two or fewer rooms. These are the units most commonly purchased for rental income—studios and small one- or two-bedroom apartments dominate the rental market in French cities. For apartments with missing room counts, those with surface area below 50 m^2 are classified as investment-type.

- **Owner-occupier-type**: Houses of any size, and apartments with three or more rooms. These are predominantly purchased for owner occupation.

This classification is imperfect—some studios are owner-occupied, some large apartments are rented—but it captures the margin along which rent control bites. Of the 621,351 transactions in the full sample, 289,545 (46.6%) are classified as investment-type and 331,806 (53.4%) as owner-occupier-type.

As a robustness check, I construct a continuous "rental score" ranging from 0 to 1, combining apartment type (+0.4), room count (+0.3 for studios/1-room, +0.15 for 2-room), and surface area (+0.2 for below 40 m^2, +0.1 for 40–70 m^2). This score provides a less discrete measure of rental market exposure.

4.4 Summary Statistics

Table 1 presents summary statistics for the full sample and key subgroups. Median transaction prices are substantially higher in treated communes (€290,000 in Paris, €230,000 in Plaine Commune) than in control cities (€173,000 in Toulouse, €155,000 in Angers). Investment-type properties command lower prices than owner-occupier properties (median €170,000 vs. €290,000), reflecting their smaller size.

Table 1: Summary Statistics

Group	N	Med. Price	Mean Price	Med. m^2	Med. €/m^2	% Apt	% Inv	Rooms
Full Sample	621,351	215,000	279,476	69	3,125	86.8	46.6	2.7
Identified Sample	451,685	199,000	233,864	72	2,727	83.4	43.8	2.8
Treated: Pre (identified)	41,134	244,000	302,389	65	3,571	89.2	40.5	2.7
Treated: Post (identified)	60,753	226,000	270,331	67	3,333	88.6	43.1	2.6
Control Cities	349,798	188,800	216,427	74	2,535	80.5	42.2	2.9
Investment-Type	289,545	170,000	252,810	40	4,545	100.0	100.0	1.6
Owner-Occupier	331,806	260,000	302,751	89	2,517	75.3	0.0	3.6

Notes: DVF transactions 2020H2–2024 in treated communes and control cities. Prices in euros. Investment-type = apartments with ≤ 2 rooms; owner-occupier-type = houses or apartments with ≥ 3 rooms. Identified sample excludes Paris and Lille (no pre-treatment data). Pre-period available only for cities adopting after 2020H2 (Plaine Commune, Est Ensemble, Lyon-Villeurbanne, Montpellier, Bordeaux).

5. Empirical Strategy

5.1 Triple-Difference Design

The primary specification is a triple-difference (DDD) model that interacts a city-level difference-in-differences with a within-city property-type comparison:

$$\ln(P_{ict}) = \alpha_c + \gamma_t + \beta_1 \text{Post}_{ct} + \beta_2 \text{Invest}_i + \beta_3(\text{Post}_{ct} \times \text{Invest}_i) + X_i'\delta + \varepsilon_{ict} \tag{5}$$

where $\ln(P_{ict})$ is the log sale price of property i in commune c at time t; α_c are commune fixed effects; γ_t are year-quarter fixed effects; $\text{Post}_{ct} = \mathbb{I}[\text{date} \geq \tau_c] \times \mathbb{I}[\text{treated}_c]$ indicates a transaction occurring in a treated commune after that commune's rent control adoption date; Invest_i indicates an investment-type property; and X_i includes controls for surface area, surface area squared, and number of rooms.

The coefficient of interest is β_3: the differential change in log prices for investment-type properties (relative to owner-occupier properties) in treated cities (relative to control cities) after rent control adoption. Under the identifying assumption that—absent rent control—the price gap between investment and owner-occupier properties would have evolved similarly in treated and control cities, β_3 estimates the causal effect of rent control on investment-type property prices.

Standard errors are clustered at the commune level throughout.

5.2 Identified Sample

The DVF five-year window creates a fundamental identification challenge: Paris (adopted July 2019) and Lille (adopted March 2020) are "always-treated" in the data—every observation for these cities is post-treatment. For these two cities, the DDD design relies entirely on control cities to identify counterfactual trends, with no ability to verify parallel pre-trends. This is problematic: Paris accounts for the majority of treated transactions and has unique market characteristics (highest rents, densest rental market in France) that make it a poor candidate for extrapolation from control cities.

I therefore define the **identified sample** as all transactions excluding Paris and Lille. This sample contains five treated city groups—Plaine Commune, Est Ensemble, Lyon-Villeurbanne, Montpellier, and Bordeaux—all of which adopted rent control between June 2021 and July 2022, providing 8–24 months of pre-treatment data in the DVF window. The identified sample has 451,685 transactions: 101,887 treated and 349,798 control. All headline results, robustness checks, and the event study use this sample. Results including Paris and Lille are reported separately as supplementary evidence.

5.3 Identification Assumptions

The DDD design requires a weaker parallel-trends assumption than a standard DiD. Instead of requiring that absolute price levels trend similarly across treated and control cities, it requires that the *relative* price of investment versus owner-occupier properties trends similarly. This is a substantive advantage: macroeconomic shocks, housing cycles, or COVID-related disruptions that affect all property types within a city are absorbed by the Post_{ct} main effect.

The assumption could be violated if, for example, a city-specific shock differentially affected small apartments relative to large ones at the same time as rent control adoption. The most plausible confound is the COVID-19 pandemic, whose lockdowns coincided with Lille's adoption (March 2020) and may have differentially affected the rental versus owner-occupier market. I address this through robustness checks excluding 2020Q3–Q4 and restricting to post-COVID adopters.

5.4 Event Study

To examine pre-trends and trace the dynamic path of the effect, I estimate an event study specification on the identified sample:

$$\ln(P_{ict}) = \alpha_c + \gamma_t + \sum_{k \neq -1} \mu_k \cdot \mathbb{I}[k_{ct} = k] \cdot \text{Treated}_c + \varepsilon_{ict} \tag{6}$$

where k indexes years relative to adoption ($k = -1$ is the reference period), estimated separately for investment-type and owner-occupier-type transactions. Relative time is binned at $[-2, +3]$ years, reflecting the available pre-treatment window (the earliest identified cohort, Plaine Commune, adopted in June 2021 with data beginning 2020H2, providing at most 11 months of pre-treatment data; later cohorts provide up to 2 years). This specification uses the `fixest` interaction-weighted estimator (Sun and Abraham, 2021) to address concerns about heterogeneous treatment effects across adoption cohorts.

5.5 City-by-City Estimation

Given concern about heterogeneity, I also estimate the full DDD specification (Equation 5, without property controls) separately for each treated city, using the full control group as comparator:

$$\ln(P_{ict}) = \alpha_c + \gamma_t + \beta_1^{(j)} \text{Post}_{ct} + \beta_2^{(j)} \text{Invest}_i + \beta_3^{(j)} (\text{Post}_{ct} \times \text{Invest}_i) + \varepsilon_{ict} \tag{7}$$

for each treated city j pooled with all 20 control cities. The coefficient of interest is $\beta_3^{(j)}$, the city-specific DDD. Results are reported for all seven cities (including Paris and Lille for completeness).

5.6 Inference Challenges

Standard inference in this setting faces two challenges. First, with 22 treated communes across 5 identified city groups and 20 control cities, the effective number of treatment clusters is small. Standard errors clustered at the commune level may understate uncertainty because communes within the same city group share a common treatment shock.

I adopt a pragmatic approach. The primary specifications cluster standard errors at the commune level (42 clusters in the identified sample: 22 treated communes and 20 control cities), which provides within-city variation in outcomes. I supplement this with two additional inference procedures: (i) a leave-one-city-out analysis that assesses the sensitivity of the pooled estimate to individual treated cities, and (ii) a randomization inference exercise that permutes treatment timing.

Second, the staggered adoption creates potential bias in two-way fixed effects (TWFE) estimators if treatment effects are heterogeneous across cohorts (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020). The DDD design partially mitigates this concern because the coefficient of interest is the triple interaction, which is less susceptible to the "negative weighting" problem that affects simple DiD (Callaway and Sant'Anna, 2021). As an explicit check, I estimate a *stacked DiD* specification that creates separate cohort-specific datasets (each treated city group paired with the full never-treated control group) and estimates the DDD on the stacked data with cohort-specific commune and time fixed effects. The stacked DDD coefficient (-0.063, $p = 0.177$) is nearly identical to the TWFE baseline (-0.055), confirming that heterogeneous treatment effects across cohorts do not materially bias the main estimates. The event study uses the interaction-weighted approach of Sun and Abraham (2021) to address treatment-effect heterogeneity.

6. Results

6.1 Main Results: Identified Sample

Table 2 presents the main results. All specifications in columns 1–6 use the identified sample (excluding Paris and Lille); column 7 reports the full-sample DDD for comparison.

Rent control does not depress property values *on average*. A simple TWFE DiD finds a negligible effect on all property types combined (-0.027, $p = 0.374$; column 1). The action is in the triple difference. Without controls, investment-type properties decline by 5.5 percent relative to owner-occupier properties in treated cities (-0.055, $p = 0.191$; column 3)—economically meaningful but statistically noisy. A stacked DiD estimator that pairs each treatment cohort with never-treated controls (addressing heterogeneous-treatment-effect bias; Goodman-Bacon 2021) yields a nearly identical -0.063 ($p = 0.177$), confirming that staggered-timing bias is negligible. Adding property-level controls for surface area, surface area squared, and room count sharpens the estimate considerably: rent control depresses investment property values by 9.3 percent (-0.093, SE $= 0.037$, $p = 0.017$; column 4). This is the headline result. The improvement with controls suggests that compositional differences between treated and control cities' property stocks add noise to the uncontrolled specification.

Column 5 uses log price per square meter as the outcome, removing the mechanical size difference between small and large properties. The DDD coefficient is -0.103 (SE $= 0.042$, $p = 0.020$)—consistent with the total-price result and suggesting that investment properties lost approximately 10 percent of their per-square-meter value relative to owner-occupier properties. Note that the Investment Type main effect flips sign in this column ($+0.269$)

relative to the total-price columns (-0.125 to -0.566): while investment-type properties have lower *total* prices (because they are smaller), they have higher *per-square-meter* prices, a standard pattern in housing markets where unit prices decrease with size.

Column 6 replaces the binary investment indicator with the continuous rental score. The rental score main effect (-1.274, $p < 0.001$) captures the large baseline price difference between high-score properties (studios, small apartments) and low-score properties (houses, large apartments)—a 1-unit shift from pure owner-occupier to pure investment type corresponds to roughly 72 percent lower prices, reflecting the enormous size and type differences along this dimension.[1] The treatment interaction is -0.155 ($p = 0.134$), negative but not statistically significant, consistent with the identified-sample result being modest. The Post $\times$ Treated main effect in this column (+0.033) has the opposite sign from the binary DDD columns; this is expected because it now represents the treatment effect at score = 0 (pure owner-occupier properties), not the average effect across all properties.

Column 7 reports the full-sample DDD including Paris and Lille: -0.230 (SE = 0.069, $p = 0.001$). For always-treated cities, $\text{Post}_{ct} = 1$ for all observations, so the "Post $\times$ Treated" coefficient for these cities is identified only from cross-sectional variation against control cities, not from before-after variation. The positive "Post $\times$ Treated" main effect in this column (+0.075) is the lower-order term in the DDD, capturing the average price level difference in treated cities relative to controls; the coefficient of interest is the triple interaction (-0.230), which measures the *differential* effect on investment-type properties. This much larger estimate is driven by the always-treated cities (particularly Paris) and should be interpreted as suggestive rather than causally identified, given the absence of pre-treatment data.

[1] Since rental score ranges from 0 (owner-occupier) to 1 (investment), the coefficient represents the full span of property types. In practice, one-standard-deviation shifts in the score are much smaller.

Table 2: Effect of Rent Control on Property Prices

	(1) TWFE	(2) TWFE+Ctrl	(3) DDD	(4) DDD+Ctrl	(5) DDD/m^2	(6) Score	(7) Full Samp.
	Identified Sample (excl. Paris & Lille)						*All Cities*
Post × Treated	−0.027	−0.008	−0.027	−0.008	−0.004	0.033	0.075*
	(0.030)	(0.029)	(0.030)	(0.029)	(0.028)	(0.063)	(0.039)
Investment Type		−0.125***	−0.566***	−0.045	0.269***		−0.566***
		(0.022)	(0.034)	(0.030)	(0.036)		(0.034)
Post × Inv.			−0.055	−0.093**	−0.103**		−0.230***
			(0.041)	(0.037)	(0.042)		(0.069)
Post × Score						−0.155	
						(0.102)	
Rental Score						−1.274***	
						(0.096)	
Controls	No	Yes	No	Yes	No	No	No
N	451,685	451,675	451,685	451,675	451,675	451,685	621,351
R^2	0.270	0.540	0.417	0.543	0.470	0.344	0.417
Commune FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Qtr FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: SE clustered at commune level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dep. var.: log price (cols. 1–4, 6–7) or log price/m^2 (col. 5). Controls: surface area, surface^2, rooms. Investment = apartments with ≤ 2 rooms. Columns 1–6 use identified sample (excl. Paris & Lille); column 7 uses full sample. Columns with controls have $N = 451{,}675$ due to 10 observations with missing surface area.

6.2 Event Study

Figure 1 presents the event study for the identified sample, estimated separately for investment-type and owner-occupier-type properties. Years relative to adoption are binned at $[-2, +3]$, with $k = -1$ as the reference period.

Two features stand out. First, the single pre-treatment coefficient ($k = -2$) is small and statistically insignificant for both property types, supporting the parallel-trends assumption: -0.039 (SE $= 0.034$, $p = 0.263$) for investment properties, and -0.014 (SE $= 0.021$, $p = 0.521$) for owner-occupier properties. Note that the $k = -2$ bin is populated exclusively by later-adopting cohorts—Montpellier and Bordeaux (both adopted July 2022), which have approximately two years of pre-treatment data—while Plaine Commune (adopted June 2021, with only 11 months of pre-treatment data) contributes only to $k = -1$ and $k = 0$. The short pre-treatment window limits the number of testable pre-periods, but the available evidence is consistent with parallel trends. Second, post-treatment coefficients are also close to zero and insignificant for both types, indicating that the modest DDD effect is diffuse rather than sharply concentrated at adoption. For investment properties, the $k = 0$ coefficient is $+0.025$ ($p = 0.632$) and $k = +2$ is -0.042 ($p = 0.518$). For owner-occupier properties, post-treatment coefficients are uniformly small and positive.

The flat event study is consistent with the modest identified-sample DDD: the effect is too small to be visually apparent in property-type-specific event studies, emerging only in the triple-difference comparison.

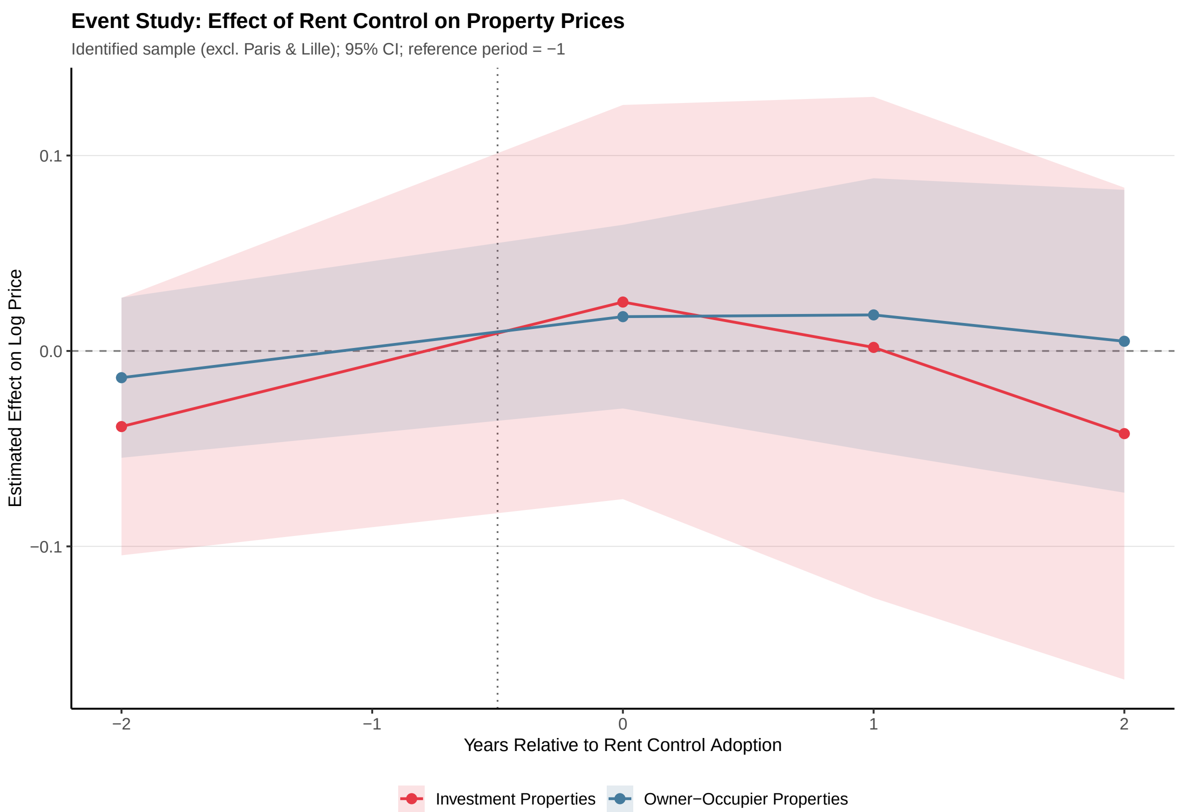

Figure 1: Event Study: Effect of Rent Control on Property Prices by Type
Notes: Identified sample (excl. Paris & Lille). Coefficients from separate regressions for investment-type and owner-occupier-type properties. Relative year $k = -1$ is the reference; $k = -2$ is populated only by Bordeaux and Montpellier (adopted July 2022, with ≈ 2 years of pre-treatment data). Earlier adopters contribute only to $k \geq -1$. Shaded bands show 95% confidence intervals. Standard errors clustered at commune level.

6.3 City-by-City Heterogeneity

Table 3 reports the DDD coefficient estimated separately for each treated city (each compared against all 20 control cities). The results are strikingly heterogeneous.

Among identified cities, Bordeaux shows a large and significant effect (-0.164, SE $= 0.037$, $p < 0.001$): investment-type properties in Bordeaux lost approximately 16 percent of their value relative to owner-occupier properties after rent control adoption. Montpellier is negative but only marginally significant (-0.073, $p = 0.062$). The remaining three identified cities—Est Ensemble, Plaine Commune, and Lyon-Villeurbanne—produce small and insignificant estimates.

In the supplementary full-sample analysis, Paris dominates with $\hat{\beta}_3 = -0.392$ (SE $= 0.057$, $p < 0.001$), suggesting a 33 percent decline in investment property values. However, this estimate cannot be causally identified because Paris has no pre-treatment data. Lille produces a positive coefficient ($+0.203$, $p < 0.001$), inconsistent with the capitalization hypothesis and possibly reflecting COVID-19 confounds (Lille adopted in March 2020, just before the first national lockdown).

Table 3: City-by-City Triple-Difference Estimates

City	DDD Coef.	Std. Error	p-value	N
Identified cities (pre-treatment data available):				
Bordeaux	-0.164^{***}	0.037	<0.001	375,354
Montpellier	-0.073^{*}	0.037	0.062	375,120
Est Ensemble	-0.039	0.061	0.531	372,739
Plaine Commune	-0.022	0.057	0.705	366,174
Lyon-Villeurbanne	$+0.006$	0.038	0.868	361,490
Always-treated cities (no pre-treatment data—supplementary):				
Paris	-0.392^{***}	0.057	<0.001	495,365
Lille	$+0.203^{***}$	0.039	<0.001	373,897

Notes: Each row reports $\hat{\beta}_3$ from the DDD specification (including all lower-order terms: Post_{ct}, Invest_i, and $\text{Post}_{ct} \times \text{Invest}_i$, with commune and year-quarter FE, without property controls) estimated on the union of the named treated city and all 20 control cities. Standard errors clustered at commune level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Paris and Lille have no pre-treatment observations in the DVF window and are reported as supplementary.

6.4 Leave-One-Out Analysis

Table 4 reports the identified-sample DDD coefficient after dropping each treated city in turn. Dropping Bordeaux—the city with the largest identified effect—reduces the coefficient from −0.055 to −0.031 ($p = 0.442$). Dropping any other city leaves the estimate between −0.057 and −0.065, all statistically insignificant. The narrow range of these four estimates is consistent with the city-by-city results: Est Ensemble, Plaine Commune, Lyon-Villeurbanne, and Montpellier each contribute small, individually insignificant effects, so removing any one of them barely shifts the pooled estimate. Bordeaux is the primary driver of the identified-sample result; the evidence for capitalization in the other four identified cities is individually weak.

Table 4: Leave-One-Out Stability of Identified-Sample DDD Estimate

Dropped City	DDD Coef.	Std. Error	p-value	N
Bordeaux	**−0.031**	**0.040**	**0.442**	**426,129**
Montpellier	−0.057	0.047	0.232	426,363
Lyon-Villeurbanne	−0.065	0.042	0.129	439,993
Plaine Commune	−0.065	0.044	0.150	435,309
Est Ensemble	−0.065	0.045	0.157	428,744

Notes: Each row reports $\hat{\beta}_3$ from the baseline DDD specification (without controls) on the identified sample after dropping the named treated city. Standard errors clustered at commune level. Bordeaux is shown in bold as the most influential city.

6.5 Size Heterogeneity

Table 5 uses a different specification to examine heterogeneity within the apartment market. Instead of the binary investment/owner-occupier classification used in the DDD, this regression restricts the sample to apartments only (excluding houses) and interacts the Post × Treated indicator with size categories (studios/1-room, 2-room as reference, 3-room, 4+ room). This within-apartment comparison does not rely on the property-type classification and provides an independent test of the capitalization mechanism: if rent ceilings bind most tightly for small units, the price decline should be strongest for studios and attenuate with size. The gradient is consistent with this prediction: studios and one-room apartments experience the largest price decline (−0.104 relative to the two-room reference category, $p = 0.001$). Three-room apartments show a positive differential (+0.054, $p = 0.043$), while four-plus-room apartments are positive but insignificant (+0.025, $p = 0.572$). This monotone gradient—from most negative for the smallest, most rental-exposed units to positive for larger units—provides direct support for the mechanism.

Table 5: Size Heterogeneity: Effect by Apartment Size Category (Identified Sample)

Post × Treated × Size	Estimate	Std. Error
× Studio/1-Room	-0.104^{***}	(0.030)
× 2-Room (ref.)	—	—
× 3-Room	$+0.054^{**}$	(0.026)
× 4+ Room	+0.025	(0.044)
Post × Treated (main)	−0.048	(0.038)
N	375,679	
R^2	0.487	

Notes: Apartments only, identified sample (N = 375,679 apartments out of 451,685 total transactions; houses excluded). Dependent variable: log price. This specification replaces the binary investment-type indicator in Table 2 with apartment size category interactions; "Post × Treated (main)" is the effect for two-room apartments (the omitted reference category), and differs from the Table 2 DDD coefficient, which pools all investment-type properties. Standard errors clustered at commune level. ** $p < 0.05$, *** $p < 0.01$.

6.6 Robustness

6.6.1 Excluding COVID

Dropping the three quarters most affected by COVID-19 (2020Q3–Q4) from the identified sample reduces the sample to 396,293 transactions. The DDD coefficient is −0.051 (SE = 0.042, $p = 0.233$), similar in magnitude to the baseline (−0.055). The result is not driven by pandemic-era distortions.

6.6.2 Post-COVID Adopters Only

Restricting to cities that adopted rent control after the pandemic (Lyon-Villeurbanne, November 2021; Montpellier, July 2022; Bordeaux, July 2022)—thereby excluding the Paris-suburban intercommunalités whose adoption overlapped with COVID-era disruptions—yields 412,368 transactions. The DDD coefficient is -0.083 (SE $= 0.050$, $p = 0.110$). The effect is negative and economically meaningful (8 percent) but statistically insignificant. Among the three post-COVID adopters, only Bordeaux shows a significant city-level estimate.

6.6.3 Randomization Inference

I conduct a randomization inference (RI) exercise on the identified sample, randomly shifting treatment adoption dates by ± 1 to 3 years and re-estimating the DDD coefficient for each of 500 permutations. For the uncontrolled DDD, the RI p-value is 0.46. For the controlled DDD (the headline specification), the RI p-value is 0.65. In both cases, the observed coefficients are not unusual under random treatment timing.

The RI result should be interpreted in context. With only five treated city groups in the identified sample and modest treatment effects, the RI has limited power. The permutation design generates variation by shifting adoption dates within the DVF window, but the window is short (4.5 years) relative to the magnitude of the shifts ($\pm$1–3 years), limiting the number of distinct treatment-control configurations. The RI does not reject the null hypothesis, but neither can it reject that a genuine modest effect exists—it is simply underpowered for this design.

6.6.4 Stacked DiD

To address the concern that TWFE estimators may produce biased estimates under staggered treatment timing with heterogeneous effects (Goodman-Bacon, 2021), I estimate a stacked DiD specification. Each cohort (treated city group) is paired with the full never-treated control group in a separate sub-dataset, and the DDD is estimated on the stacked data with cohort-specific commune and time fixed effects. The stacked DDD coefficient without controls is -0.063 (SE $= 0.046$, $p = 0.177$), virtually identical to the TWFE baseline of -0.055 ($p = 0.191$). With controls, the stacked DDD yields -0.100 (SE $= 0.040$, $p = 0.016$), confirming that the controlled headline result (-0.093) is robust to staggered-timing bias.

6.6.5 Sensitivity to Control Variables

A potential concern is that the shift from the uncontrolled DDD (-0.055, $p = 0.191$) to the controlled DDD (-0.093, $p = 0.017$) is driven by room-count controls that overlap with the investment-type classification. I address this by estimating a specification controlling only for surface area and surface area squared, without room count. The resulting DDD coefficient is -0.094 (SE $= 0.037$, $p = 0.015$), nearly identical to the full-controls estimate. The improvement with controls therefore reflects absorption of within-type size variation (small studios versus large studios) rather than mechanical reclassification along the investment/owner-occupier margin.

6.6.6 Event Study Pre-Trends

As reported above, the event study shows no evidence of differential pre-trends. The $k = -2$ coefficient for investment properties is -0.039 ($p = 0.263$) and for owner-occupier properties is -0.014 ($p = 0.521$). While the pre-treatment window is short (only 1–2 years for most cities), the absence of significant pre-trends supports the identifying assumption.

6.7 Transaction Composition

If rent control depresses the value of investment properties, it may also affect the composition of transactions. Figure 6 plots the share of investment-type transactions over time for treated and control cities. Treated cities consistently show a higher investment-type share (reflecting their larger rental markets), but the two groups move in parallel over time. Crucially, there is no sharp compositional break at the time of adoption, suggesting that the DDD results are not driven by a sudden change in which types of properties are transacted.

6.8 Summary of Findings

The empirical evidence supports a cautious conclusion. The identified-sample DDD estimate is modest: -0.055 without controls ($p = 0.191$) and -0.093 with controls ($p = 0.017$). The controlled estimate is robust to dropping room-count controls (-0.094 with surface only), to stacked DiD estimation (-0.100), and to excluding COVID-affected quarters (-0.051). However, randomization inference does not reject the null ($p = 0.46$ uncontrolled, $p = 0.65$ controlled), reflecting the limited power of five treated city groups. The effect is concentrated in Bordeaux (-0.164, $p < 0.001$) and absent in three of the five identified cities. The within-apartment size gradient provides the strongest support for the capitalization mechanism: studios and one-room apartments lose value relative to larger units in treated cities. In the supplementary full-sample analysis, Paris shows a large suggestive effect (-0.392), consistent with severe bindingness in France's tightest market, but this cannot be causally identified.

7. Discussion

7.1 Why Bordeaux?

Among the five identified cities, Bordeaux stands out as the only one with a large, statistically significant capitalization effect (-0.164). Several features distinguish Bordeaux. First, it has the highest share of small rental apartments among non-Parisian adopters, with 43.3 percent of transactions classified as investment-type. Second, Bordeaux experienced a decade-long property boom prior to 2022, with prices roughly doubling between 2012 and 2022, suggesting a market where rents may have risen above the regulatory ceiling. Third, Bordeaux is a standalone commune (not an intercommunalité like Plaine Commune or Est Ensemble), so the regulation applies uniformly across the city, potentially reducing measurement error in treatment assignment.

The null results for Plaine Commune, Est Ensemble, and Lyon-Villeurbanne are also informative. These are markets where either rents were already near the ceiling (so the regulation does not bind) or where the intercommunalité structure creates heterogeneity in treatment intensity across communes. The capitalization mechanism requires the regulation to actually constrain rents—where it does not, no price effect is expected.

7.2 The Paris Question

Paris looms large in the data but cannot be causally identified. The city-specific DDD of -0.392 is striking and consistent with the capitalization model: Paris has the tightest rental market in France, the widest gap between market rents and reference rents, and the highest share of investment-type transactions (60.4%). A back-of-envelope calculation illustrates the magnitudes: a studio in the 11th arrondissement renting at €1,200/month, where the *loyer majoré* is €900/month, faces an annual rent reduction of €3,600, which capitalized at 4 percent implies a property value decline of €90,000—approximately 25 percent for a typical €350,000 studio.

However, without pre-treatment data, this estimate may reflect pre-existing trends, COVID-era disruptions, Airbnb regulation effects, or other confounds. The Paris estimate should be viewed as an upper bound on the capitalization effect in a severely binding regulatory environment, not as a causally identified treatment effect.

7.3 Mechanisms and Alternative Explanations

The capitalization interpretation requires that the effect operates through changes in expected rental income. Several alternative mechanisms deserve consideration.

Composition effects. If rent control induced a shift in the type of investment properties transacted—for example, by discouraging sales of high-quality studios and encouraging sales of lower-quality units—the observed price decline could reflect compositional change rather than true depreciation. Reassuringly, adding property controls (surface area, rooms) strengthens rather than weakens the DDD coefficient (from -0.055 to -0.093), suggesting that composition changes do not drive the result. If anything, controlling for observable property characteristics reveals a cleaner capitalization signal. However, unobservable quality differences could still play a role.

Demand reallocation. If rent control made the rental market less attractive, some potential landlords may have shifted toward owner-occupation of small apartments, while some potential tenants may have sought to purchase rather than rent. This could depress demand (and prices) for investment-type properties while increasing demand for owner-occupier properties. Both channels are consistent with the regulation having real economic effects; they differ only in the precise mechanism.

Investor exit. If the *encadrement* signaled increasingly regulated markets, forward-looking investors may have sold properties preemptively, increasing supply and depressing prices. This channel is observationally equivalent to capitalization but operates through the supply side of the sales market.

Confounding city-specific trends. The most serious concern is that Bordeaux experienced a differential trend in investment-to-owner-occupier price ratios unrelated to rent control. Bordeaux-specific developments—the arrival of the LGV high-speed rail in 2017, changing demographics, or tourism-driven short-term rental regulation—could drive such a trend. The clean pre-trends in the event study provide some reassurance, but the pre-treatment window is short (approximately 2 years for Bordeaux).

7.4 Limitations

Several limitations temper the conclusions.

Modest effect size and borderline significance. The identified-sample DDD without controls is not statistically significant ($p = 0.191$). Significance requires property-level controls ($p = 0.017$). This sensitivity to specification, while common in applied work, suggests that the effect is modest in magnitude and at the edge of what the data can reliably detect.

Concentration in one city. The leave-one-out analysis shows that dropping Bordeaux renders the identified-sample DDD insignificant. While Bordeaux represents a genuine natural experiment with clean pre-treatment data, the concentration of the result in a single city limits external validity.

Measurement error in rental exposure. The classification of investment-type versus owner-occupier-type properties is based on observables (property type, room count, surface area) rather than actual tenure status. Classical measurement error in the treatment-exposure variable would attenuate the DDD coefficient toward zero, meaning the true effect on genuinely rented properties may be larger. The continuous rental score and size gradient specifications partially address this.

Non-significant randomization inference. The RI yields $p = 0.46$, which would normally undermine confidence in the result. However, with only five treated city groups and a short data window, the RI has limited power. The test is informative about whether the effect is unusually large given random timing, but cannot rule out a genuine modest effect that the permutation test is underpowered to detect.

Short pre-treatment window. The identified cities have at most 2 years of pre-treatment data, limiting the ability to verify long-run parallel trends. Because the DVF is a five-year rolling window whose start date advances with each release, future releases will not extend pre-treatment coverage for these cities—the earliest observations will shift forward rather than backward. Extending the pre-treatment window would require access to archived historical DVF extracts or alternative transaction data sources.

Non-random adoption. The staggered adoption of the *encadrement des loyers* was driven by political choice, not random assignment. Cities that adopted had tighter housing markets and stronger political support for regulation. If these same characteristics independently drove differential trends in investment versus owner-occupier prices, the DDD estimate would be biased. The control cities were selected to be comparable large cities ($>$ 100,000 population), but selection on unobservables cannot be ruled out.

General equilibrium effects. If rent control in treated cities caused investors to redirect capital to owner-occupier properties (driving up their prices) or to control cities (depressing returns there), the third difference would overstate the capitalization effect. Similarly, the stable unit treatment value assumption may be violated if treated and control cities interact through housing markets. Investors displaced from Bordeaux by rent control may purchase properties in Toulouse or Nantes (control cities), affecting prices there.

7.5 Magnitudes and Welfare

A rough welfare calculation helps quantify the stakes for Bordeaux, the city with the strongest identified effect. Bordeaux had approximately 25,000 property transactions in the DVF sample over 4.5 years, of which 43 percent are investment-type. If the DDD estimate of -0.164 log points represents genuine capitalization, and the median investment-type property in Bordeaux sells for approximately €185,000, the implied per-transaction wealth transfer is approximately €28,000 ($= 185{,}000 \times (1 - e^{-0.164})$). Over approximately 2,400 investment-type transactions per year, this implies an aggregate wealth redistribution of approximately €67 million per year in Bordeaux.

For Paris—if the suggestive full-sample estimate is taken at face value—the magnitudes are an order of magnitude larger. With approximately 19,000 investment-type transactions per year and a median price of €300,000, the DDD coefficient of -0.392 implies a per-transaction wealth transfer of approximately €97,500 and an aggregate of approximately €1.9 billion per year. These calculations are necessarily crude: they assume the entire DDD coefficient represents pure capitalization, apply the median price uniformly, and ignore that many transactions would occur at similar prices absent rent control. They are intended to illustrate the potential economic magnitude, not to provide definitive welfare estimates.

On the landlord side, the wealth loss from capitalization may be partially offset by tax benefits. French landlords can deduct rental losses from taxable income under certain conditions (*déficit foncier*), and lower property values reduce the *taxe foncière* (property tax) and *Impôt sur la Fortune Immobilière* (IFI, the wealth tax on real estate above €1.3 million). These offsets are likely small relative to the capitalization loss but deserve consideration in a complete welfare analysis.

8. Conclusion

France's staggered adoption of the *encadrement des loyers* provides a natural experiment to test whether rent control is capitalized into property sale prices. Using 451,685 property transactions from the identified sample—five cities with pre-treatment data—and a triple-difference design comparing investment-type to owner-occupier properties, I find evidence of modest capitalization. The DDD estimate with controls is -0.093 ($p = 0.017$), suggesting a 9 percent price decline for investment-type properties relative to owner-occupier properties in treated cities. The within-apartment size gradient confirms the mechanism: studios and one-room apartments show the strongest effects.

But this finding comes with important qualifications. The baseline specification without controls is not statistically significant. The effect is concentrated in Bordeaux, with null results in four of five identified cities. And the randomization inference, while underpowered, does not reject the null. The honest conclusion is that rent control may depress investment property values where the regulation binds severely—Bordeaux, and suggestively Paris—but the effect is too modest and localized to constitute strong evidence for a general capitalization channel.

The size gradient provides the most compelling evidence for the mechanism. Within the apartment market, the monotone pattern—from negative for studios to positive for three-room apartments—is precisely what the capitalization model predicts and difficult to explain by alternative mechanisms. This within-market comparison is the strongest test of whether rent ceilings on small rental units translate into differential asset price effects.

The policy implications are nuanced. For cities where rent control meaningfully binds—Bordeaux and, suggestively, Paris—the results suggest that the regulation achieves a form of wealth redistribution from owners of small rental apartments to the broader market, with economically significant magnitudes. But extending rent control to cities where market rents are near or below the ceiling is unlikely to produce property value effects. The expansion of the *encadrement* to additional French cities may prove largely symbolic in asset-price terms.

Several avenues for future research emerge. First, access to archived historical DVF extracts—or alternative transaction data sources such as notarial databases—would enable observation of pre-treatment periods for Paris (adopted 2019) and Lille (adopted 2020), allowing verification of whether the large suggestive effects documented here reflect genuine capitalization or pre-existing trends. Because the DVF is a five-year rolling window, standard public releases will not extend backward to cover these cities' pre-treatment periods. Second, linking DVF transaction data to rental market data—such as the *observatoire des loyers* reference rents—would enable direct measurement of the regulatory bite for individual properties, allowing a more precise estimate of the capitalization elasticity. Third, the ongoing expansion of rent control to new French cities (Grenoble adopted in 2024) creates additional natural experiments.

More broadly, this paper demonstrates that rent control's effects extend beyond the rental market. The standard policy debate focuses on tenants, landlords, and housing supply. The capitalization channel documented here adds a fourth dimension: the effect on property wealth. A complete evaluation of rent control policy must account for these asset-price consequences alongside the more commonly studied rental-market effects. In the debate over housing affordability, rent control is often treated as a transfer from landlord to tenant; these results suggest it is also a permanent revaluation of urban wealth.

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

Adda, Jérôme and Sara Pinoli, "Rent Control and Misallocation," Technical Report 18012, CEPR Discussion Paper 2023.

Ahlfeldt, Gabriel and Wolfgang Maennig, "After the Mietendeckel: The Effect of Rent Control on Rents and Property Values in Berlin," Technical Report 16894, CEPR Discussion Paper 2022.

Arnott, Richard, "Tenancy Rent Control," *Swedish Economic Policy Review*, 1995, *2* (1), 89–121.

_ , "Time for Revisionism on Rent Control?," in "Journal of Economic Perspectives," Vol. 9 1995, pp. 99–120.

Autor, David H, Christopher J Palmer, and Parag A Pathak, "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts," *Journal of Political Economy*, 2014, *122* (3), 661–717.

Bonnefoy, Mathilde, Guillaume Chapelle, and Marianne Sagot, "L'encadrement des loyers à Paris: premiers effets sur les loyers et les marchés," *Revue d'économie régionale et urbaine*, 2021, *4*, 697–730.

Bono, Pierre-Henri and Alain Trannoy, "Houses and Apartments in French Urban Markets: What Role for Supply Constraints?," *Regional Science and Urban Economics*, 2022, *97*, 103817.

Breidenbach, Philipp, Konstantin A Kholodilin, and Andreas Mense, "Rent Control and Property Prices: Not a Zero-Sum Game," Discussion Paper 2011, DIW Berlin 2022.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Chapelle, Guillaume and Benjamin Vignolles, "Housing Market and Urban Planning Regulations: Evidence from France," *Annals of Economics and Statistics*, 2019, (136), 65–100.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Diamond, Rebecca, Tim McQuade, and Franklin Qian, "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco," *American Economic Review*, 2019, *109* (9), 3365–3394.

Early, Dirk W, "Rent Control, Reputation and Quality Maintenance," *Journal of Real Estate Finance and Economics*, 2000, *20*, 117–133.

Glaeser, Edward L and Erzo FP Luttmer, "Does Rent Control Reduce Segregation?," *Swedish Economic Policy Review*, 2003, *10*, 179–202.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Gyourko, Joseph and Peter Linneman, "Rent Controls and Rental Housing Quality: A Note on the Effects of New York City's Old Controls," *Journal of Urban Economics*, 1990, *27* (3), 398–409.

Harding, John P, Stuart S Rosenthal, and C F Sirmans, "The Capitalization of Income Streams and the Valuation of Closed-End Funds: Taxes, Losses and Discounts," *Review of Economics and Statistics*, 2012, *95* (5), 1617–1631.

Mense, Andreas, Claus Michelsen, and Konstantin A Kholodilin, "Secondary Effects of Rent Control on Property Values," *Journal of Urban Economics*, 2023, *133*, 103530.

Olsen, Edgar O, "An Econometric Analysis of Rent Control," *Journal of Political Economy*, 1972, *80* (6), 1081–1100.

Sims, David P, "Out of Control: What Can We Learn from the End of Massachusetts Rent Control?," *Journal of Urban Economics*, 2007, *61* (1), 129–151.

Stull, William J, "The Implicit Tax on Rent-Controlled Tenants," *American Economic Review*, 1978, *68* (5), 930–941.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

A. Data Appendix

A.1 DVF Data Source

The *Demandes de Valeurs Foncières* (DVF) is published by the Direction Générale des Finances Publiques (DGFiP) and made available through `data.gouv.fr`. The dataset records every notarized property transaction in metropolitan France, extracted from the land registry (*cadastre*). The data are available as a five-year rolling window; as of the data collection date (March 2026), the available window covers transactions from 2020H2 through 2024.

File format: DVF files are pipe-delimited text files with French decimal comma formatting. Each mutation (transaction) may span multiple rows corresponding to different parcels, lots, or buildings within the transaction.

Key variables:

- `valeur_fonciere`: Transaction price in euros.
- `type_local`: Property type (Maison, Appartement, Local industriel, Dépendance).
- `surface_reelle_bati`: Building surface area in m^2.
- `nombre_pieces_principales`: Number of main rooms.
- `code_departement`, `code_commune`: Geographic identifiers. The 5-digit INSEE commune code is constructed as `sprintf("%02d%03d", departement, commune)`.
- `date_mutation`: Transaction date (DD/MM/YYYY format).

Exclusions: Départements 57 (Moselle), 67 (Bas-Rhin), and 68 (Haut-Rhin) are excluded from DVF due to the Alsace-Moselle land registry system (*livre foncier*). Overseas territories are excluded by restricting to 5-digit commune codes.

A.2 Sample Filters

The following filters were applied sequentially:

1. Residential sales only: `nature_mutation = "Vente"` and `type_local` $\in$ `{"Maison", "Appartement"}`.
2. Non-missing, positive transaction price.
3. Exclude Alsace-Moselle (départements 57, 67, 68).

4. Price between €10,000 and €50,000,000.

5. Surface area between 5 and 5,000 m^2 (if available).

6. Price per m^2 between €200 and €30,000 (if computable).

7. Metropolitan France only (5-digit commune codes).

After all filters: 5,297,214 transactions. The full analysis sample (treated communes ∪ control cities): 621,351. The identified sample (excluding Paris and Lille): 451,685.

A.3 Treatment Assignment

Treatment is assigned at the commune level using the 5-digit INSEE code. Table 6 lists all 45 treated communes, their city group, and adoption date.

Table 6: Treatment Assignment: Communes and Adoption Dates

City Group	Communes (INSEE codes)	Adoption Date	N
Paris†	75101–75120 (20 arr.)	Jul. 1, 2019	145,567
Lille†	59350, 59298, 59360	Mar. 1, 2020	24,099
Plaine Commune	93001, 93007, 93027, 93031, 93039, 93047, 93066, 93070, 93079	Jun. 1, 2021	16,376
Lyon-Villeurbanne	69123, 69266	Nov. 1, 2021	11,692
Est Ensemble	93005, 93006, 93010, 93048, 93049, 93053, 93055, 93061, 93064	Dec. 1, 2021	22,941
Montpellier	34172	Jul. 1, 2022	25,322
Bordeaux	33063	Jul. 15, 2022	25,556
Total treated transactions			**271,553**
Identified sample treated transactions			**101,887**

† Always-treated in DVF window (no pre-treatment observations). Excluded from identified sample.

A.4 Control Cities

The control group consists of 20 French cities with populations above 100,000 that did not adopt the *encadrement des loyers* during the sample period. Control city selection targeted large, urbanized communes with active housing markets. Two initially selected cities

(Strasbourg, code 67482; Metz, code 57463) were automatically dropped because they are in the excluded Alsace-Moselle départements.

Control cities: Toulouse (31555), Nantes (44109), Nice (06088), Rennes (35238), Rouen (76540), Toulon (83137), Saint-Étienne (42218), Le Havre (76351), Reims (51454), Dijon (21231), Angers (49007), Clermont-Ferrand (63113), Tours (37261), Limoges (87085), Amiens (80021), Perpignan (66136), Brest (29019), Besançon (25056), Orléans (45234), Caen (14118).

B. Identification Appendix

B.1 Pre-Treatment Data Availability

Table 7 summarizes data availability by city group.

Table 7: Pre-Treatment Data Availability

City	Adoption Date	Pre-Treatment Obs.	Post-Treatment Obs.
Paris†	July 2019	0	145,567
Lille†	March 2020	0	24,099
Plaine Commune	June 2021	3,612	12,764
Est Ensemble	December 2021	7,760	15,181
Lyon-Villeurbanne	November 2021	4,196	7,496
Montpellier	July 2022	12,285	13,037
Bordeaux	July 2022	13,281	12,275

† No pre-treatment data in DVF window. Excluded from identified sample.

Paris and Lille have no pre-treatment data in the DVF window. All headline results use the identified sample (five cities with pre-treatment data). Paris and Lille are reported in city-by-city analysis as supplementary evidence.

B.2 Randomization Inference Details

The RI procedure randomly shifts treatment adoption dates for identified-sample treated communes by a uniform draw from $[-1095, -365] \cup [365, 1095]$ days (1–3 years in either direction). For each of 500 permutations, I re-compute Post_{ct} and re-estimate the DDD coefficient on the identified sample.

The resulting permutation distribution has mean -0.040 and standard deviation 0.036. The actual estimate (-0.055) falls within one standard deviation of the permutation mean, yielding a two-sided RI p-value of 0.46. With only five treated city groups and a short data window, the permutation test has limited power to distinguish a genuine modest effect from random variation.

C. Robustness Appendix

C.1 Excluding COVID-Affected Quarters

Dropping 2020Q3–2020Q4 from the identified sample removes approximately 55,000 transactions (12% of the sample). The DDD coefficient is -0.051 (SE $= 0.042$, $p = 0.233$), virtually identical to the baseline -0.055.

C.2 Post-COVID Adopters Only

Restricting to Lyon-Villeurbanne, Montpellier, and Bordeaux (adopted November 2021–July 2022) with all 20 control cities yields 412,368 transactions. The DDD coefficient is -0.083 (SE $= 0.050$, $p = 0.110$). The negative point estimate is consistent with the capitalization mechanism but statistically insignificant, reflecting lower bindingness and smaller sample sizes.

D. Additional Figures

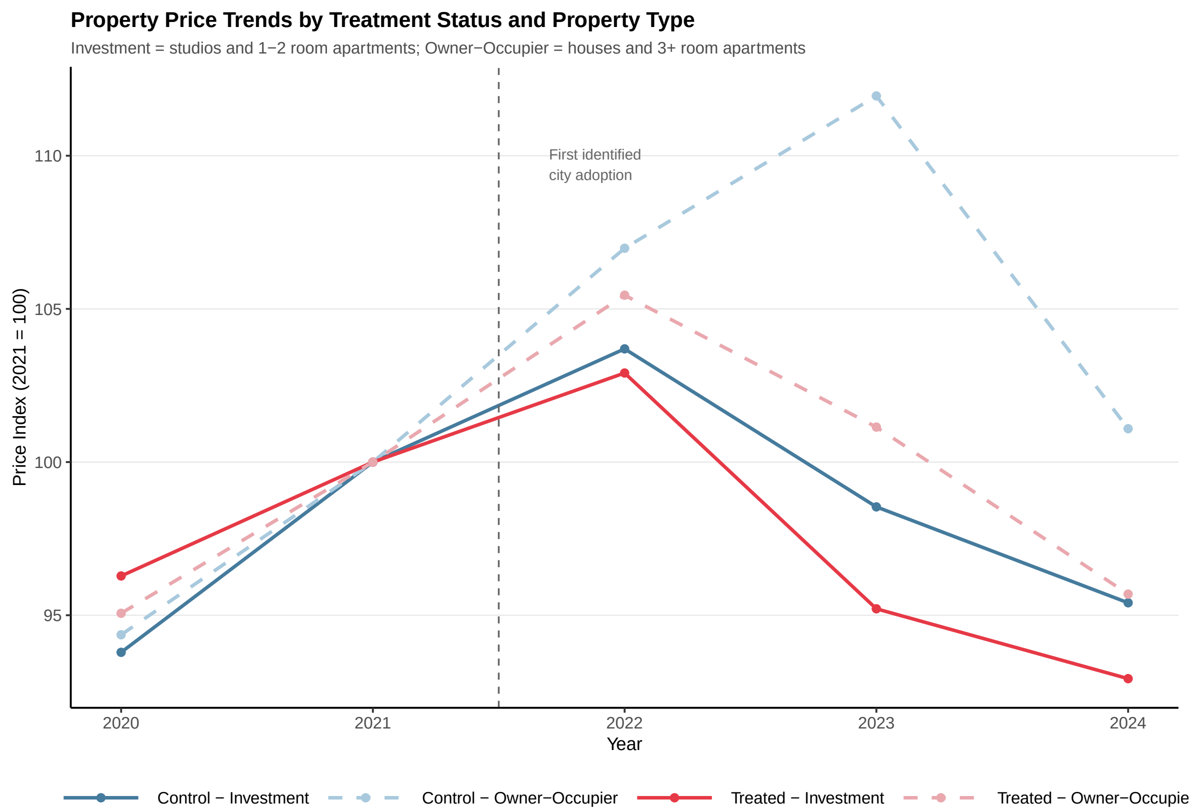

Figure 2: Property Price Trends by Treatment Status and Property Type
Notes: Price index normalized to 2021 = 100. "Investment" = studios and 1–2 room apartments. "Owner-Occupier" = houses and 3+ room apartments. Dashed vertical line marks the first identified-sample city adoption (Plaine Commune, June 2021). Source: DVF 2020H2–2024.

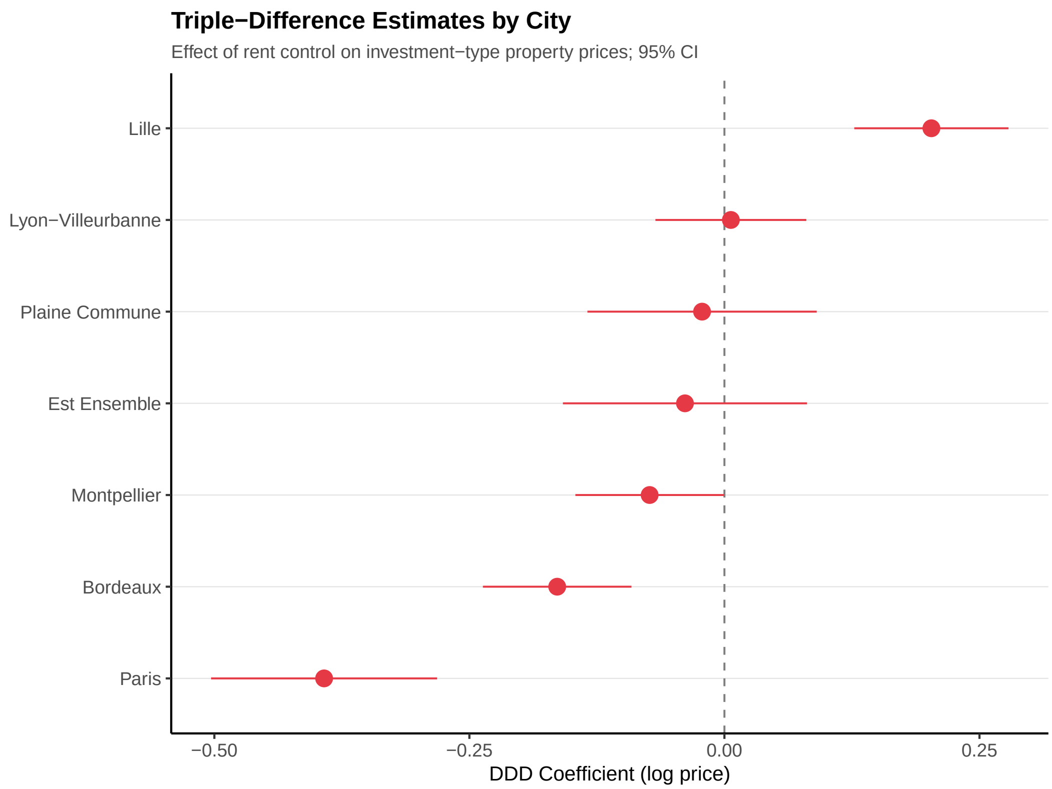

Figure 3: Triple-Difference Estimates by City

Notes: Point estimates and 95% confidence intervals for city-specific DDD coefficients. Each city is compared against the full set of 20 control cities. Standard errors clustered at commune level.

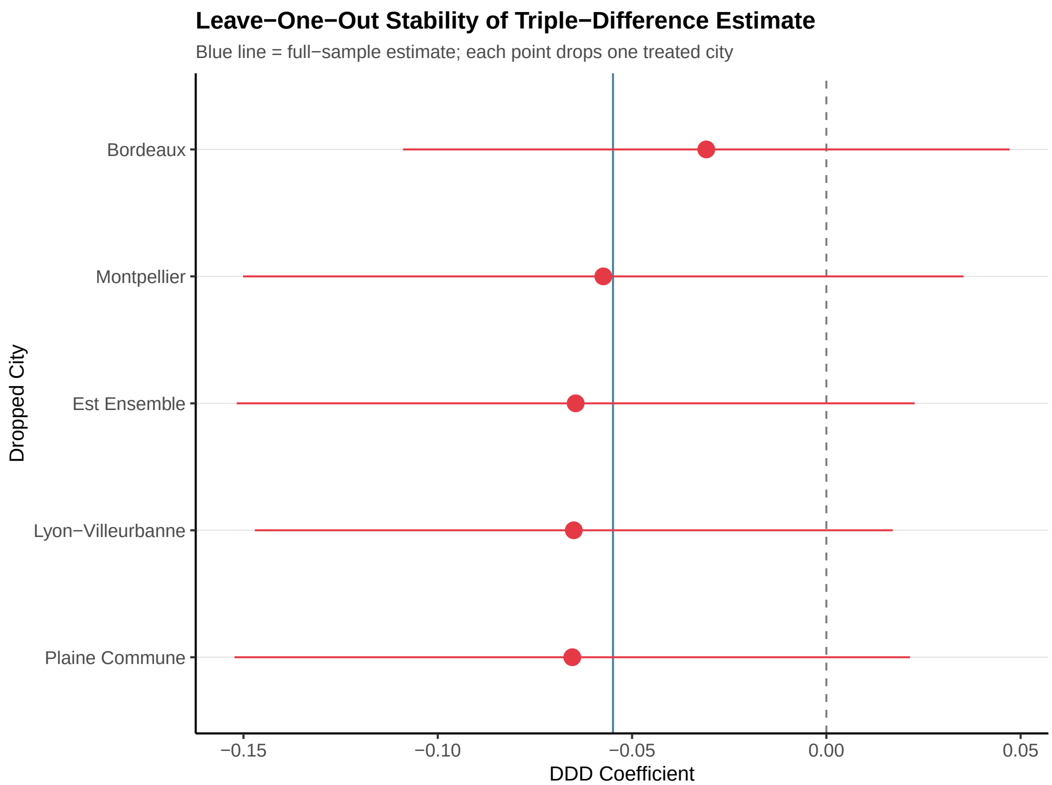

Figure 4: Leave-One-Out Stability of the Identified-Sample DDD Estimate
Notes: Each point shows the identified-sample DDD coefficient after dropping the named treated city. Blue horizontal line = full identified-sample estimate. Bars show 95% confidence intervals.

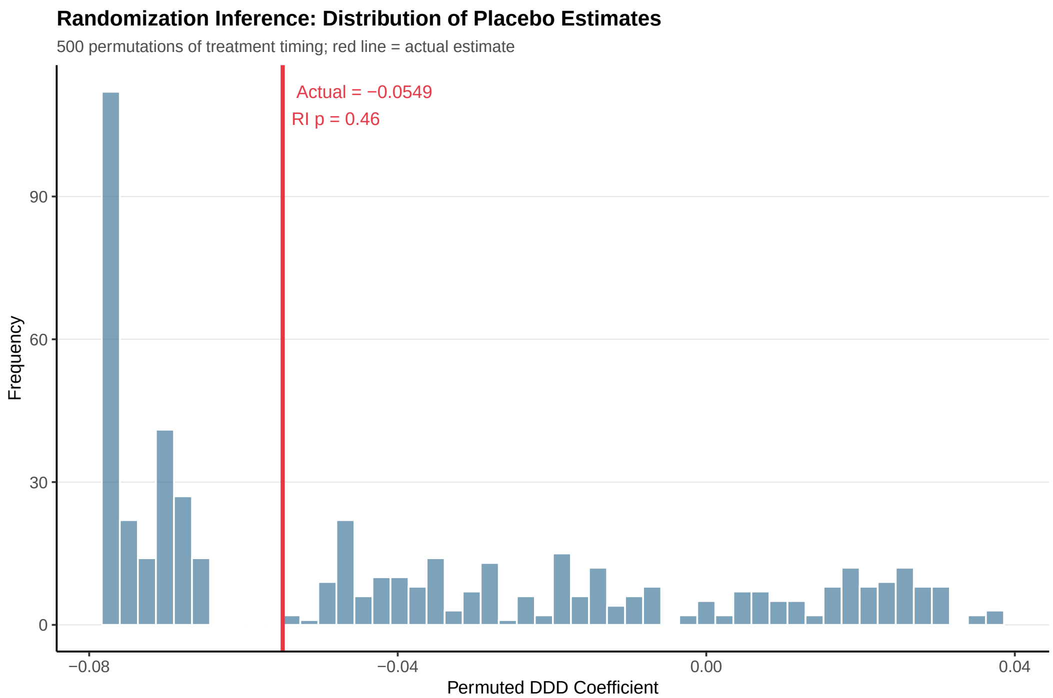

Figure 5: Randomization Inference: Distribution of Placebo DDD Coefficients
Notes: Distribution of DDD coefficients from 500 permutations of treatment timing (±1–3 years) on the identified sample. Red vertical line = actual estimate (−0.055). RI p-value = 0.46.

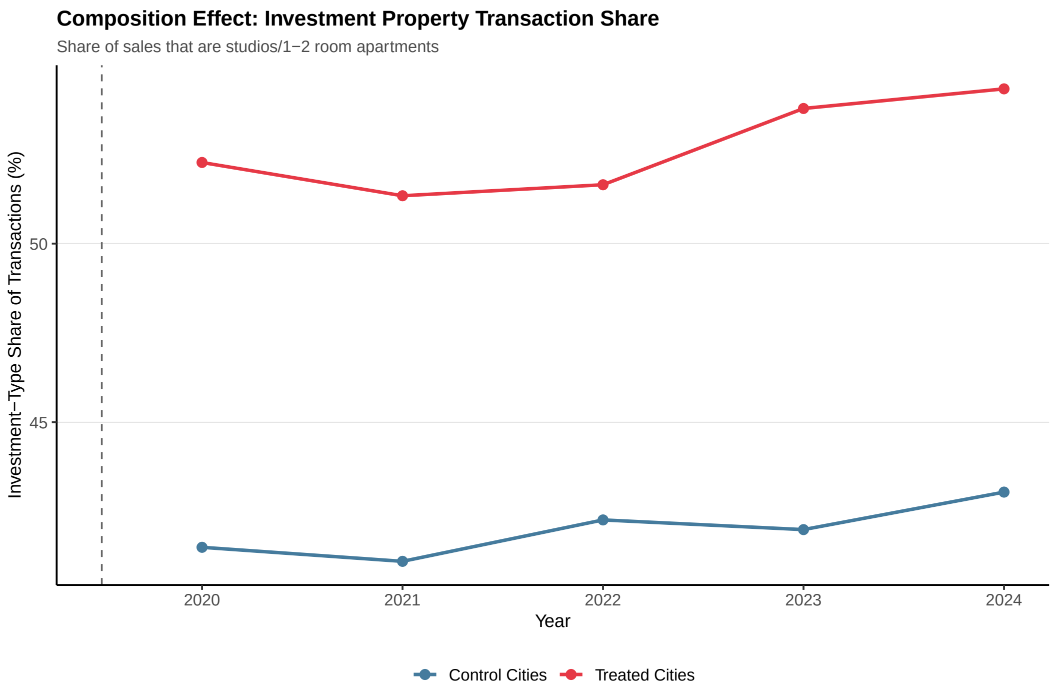

Figure 6: Investment-Type Transaction Share Over Time

Notes: Share of transactions classified as investment-type (studios and 1–2 room apartments) by year and treatment status. Identified sample ($N = 451{,}685$).