

The Subsidy Cap Trap: Developer Pricing at Government Ceilings in the French Housing Market

APEP Autonomous Research* @ai1scl

March 30, 2026

Abstract

Housing subsidies with price ceilings create incentives for developers to price at the cap, capturing the subsidy as profit rather than passing it through to buyers. Using the universe of 2.9 million French property transactions and the July 2024 reclassification of 865 communes into higher PTZ (Prêt à Taux Zéro) zones, I estimate that new-build transactions exhibit 43% excess mass at PTZ price caps—six times larger than resale transactions at the same price points. When reclassification shifted caps by €22,500–€52,500, bunching at old caps declined significantly for treated communes (triple-difference: -0.124 , $p = 0.028$). Placebo caps 10–20 thousand euros from the true threshold show no excess mass. These findings reveal a *subsidy cap trap*: well-intentioned price ceilings become focal points for developer pricing, limiting the buyer-side benefits of subsidized lending.

JEL Codes: R31, H24, R38

Keywords: bunching, housing subsidies, price ceilings, developer capture, PTZ, France

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 23m).

1. Introduction

When governments subsidize homeownership with price ceilings, who actually benefits? The intended logic is straightforward: cap the purchase price, provide favorable financing below the cap, and help households buy homes they could not otherwise afford. But developers face the same cap. If they can set prices at or just below the ceiling—capturing the maximum revenue while keeping buyers eligible for the subsidy—the price ceiling becomes a price floor. The subsidy flows to producers, not consumers.

This paper provides the first evidence of developer pricing at government-imposed subsidy ceilings in the French housing market. France’s Prêt à Taux Zéro (PTZ)—a zero-interest loan program for first-time buyers of new-build housing—imposes 20 distinct price caps that vary by geographic zone and household size. Using the universe of French property transactions from 2022 to 2024, I document sharp bunching of new-build sales at PTZ caps: 43% excess mass at the zone B2 cap of €165,000, six times larger than the 7% excess mass observed for resale transactions at the same price point. Resale transactions, which are ineligible for PTZ, serve as a clean placebo: the absence of economically meaningful bunching confirms that the pattern is driven by the subsidy mechanism, not by general features of the price distribution.

The July 2024 reclassification of 865 communes into higher PTZ zones provides a powerful causal test. The *Arrêté* of 5 July 2024 shifted applicable price caps by €22,500 to €52,500 for reclassified communes, while leaving neighboring communes in the same old zone untouched. If bunching reflects developer pricing at subsidy limits, it should *migrate*: excess mass at the old cap should disappear in reclassified communes and reappear at the new cap. I estimate a triple-difference—near-cap \times reclassified \times post—that yields a coefficient of -0.124 ($p = 0.028$), confirming that bunching at the old cap declined differentially for treated communes after reclassification.

These results speak to a first-order question in housing policy: the incidence of buyer subsidies. A large literature documents that housing vouchers and tax credits are partially capitalized into prices (Susin, 2002; Fack, 2006; Eriksen and Rosenthal, 2015), and that mortgage subsidies inflate house prices when supply is constrained (Glaeser et al., 2008; Mian and Sufi, 2009). But the mechanism through which price ceilings on subsidized purchases translate into developer pricing behavior has received little attention. The only prior study of France’s PTZ, Gobillon and le Blanc (2008), examined homeownership effects using survey data and did not analyze pricing responses or bunching at subsidy caps.

This paper contributes to three literatures. First, it extends the bunching methodology (Saez, 2010; Chetty et al., 2011; Kleven, 2016) to housing subsidy caps, showing that price ceilings on subsidized purchases generate bunching analogous to tax notches (Kleven and

Waseem, 2013; Slemrod, 2013). Unlike tax notches, where the individual faces the threshold, here a third party—the developer—responds to the buyer’s subsidy eligibility. This is closer to the employer-side bunching documented in labor markets, but in a housing context where the “firm” is a developer and the “kink” is a price ceiling.

Second, it provides a novel identification test within the bunching paradigm: *difference-in-bunching* from policy reclassification. When administrative boundaries shift, the bunching mass should follow. This migration test is stronger than static bunching estimation because it eliminates time-invariant confounders (e.g., round-number preferences, construction cost clustering) that could mimic bunching at price caps. The idea complements the “bunching at notches” literature (Best and Kleven, 2018; Kopczuk and Munroe, 2015) by exploiting a *moving* threshold rather than a static one.

Third, it documents a specific mechanism—the *subsidy cap trap*—through which well-intentioned housing policies fail to benefit their intended recipients. When developers price at the cap, the PTZ subsidy is fully absorbed into revenue: the buyer pays the maximum eligible price and receives the zero-interest loan, but the effective house price is set by the developer’s strategic response to the cap, not by market fundamentals. This mechanism is distinct from the general supply-side capitalization documented by Hilber and Vermeulen (2016) and Diamond and McQuade (2019), because it operates through a discrete price ceiling rather than through aggregate demand shifts.

The remainder of the paper is organized as follows. Section 2 describes the PTZ program and the 2024 reclassification. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

The PTZ program. France’s Prêt à Taux Zéro (PTZ), created in 1995 and reformed multiple times, provides zero-interest loans to first-time buyers purchasing new-build housing. The loan covers a fraction of the purchase price, with the fraction and repayment schedule determined by household income, zone, and household size. Critically, the PTZ imposes a *plafond d’opération*—a maximum eligible purchase price—that depends on two variables: the commune’s housing-tension zone and the buyer’s household size.

Five zones classify France’s 34,875 communes by housing market pressure: Abis (Île-de-France core), A (other major agglomerations), B1 (secondary cities and their peripheries), B2 (smaller towns), and C (rural areas). Combined with four household-size brackets (1, 2, 3, and 4+ persons), this creates 20 distinct price caps. For a two-person household, the caps range from €150,000 in zone C to €225,000 in zones A and Abis. In zone B2, the cap is

€165,000; in B1, €202,500. These caps create a discrete set of focal points in the transaction price distribution.

The July 2024 reclassification. The *Arrêté* of 5 July 2024 reclassified 865 communes into higher housing-tension zones. The largest group—485 communes—moved from B2 to B1, shifting the applicable two-person-household cap from €165,000 to €202,500, a €37,500 increase. An additional 190 communes moved from C to B1 (+€52,500), 129 from B1 to A (+€22,500), and 42 from A to Abis. Approximately 28,000 communes with B2 or lower classifications remained unchanged, providing a natural control group.

The reclassification was driven by updated housing market indicators (price levels, construction activity, demographic pressure) and applied uniformly to qualifying communes. Individual developers had no influence over which communes were reclassified. The effective date was immediate upon publication (July 11, 2024), giving developers little time to anticipate the change.

Developer incentives. A developer selling a new-build unit faces a trade-off. Pricing above the PTZ cap means the buyer loses subsidy eligibility, potentially reducing demand. Pricing below the cap means leaving revenue on the table. The profit-maximizing response is to price at or just below the cap, capturing the maximum revenue while keeping the unit PTZ-eligible. This creates a “subsidy cap trap”: the cap intended to protect buyers becomes the developer’s target price.

3. Data

DVF transactions. I use the universe of French property transactions from the *Demandes de Valeurs Foncières* (DVF), published as open data on data.gouv.fr. DVF records every property sale in metropolitan France with the transaction price, commune code, date, property type, and surface area. I use transactions from 2022 through 2024, yielding 2,916,013 observations across 34,875 communes after zone matching (93.5% match rate).

New-build identification. I identify new-build transactions through the *Vente en l’État Futur d’Achèvement* (VEFA) designation, which indicates off-plan sales—the standard mechanism for new-build housing in France. VEFA transactions account for 2.2% of all sales (64,562 transactions), consistent with the national share of new-build activity. These are the transactions eligible for PTZ financing.

Zone classification and reclassification. I use the official commune-level zone classification from the Ministry of Housing (data.gouv.fr) and the published list of 865 reclassified communes

from the July 2024 *Arrêté*. I assign each transaction its applicable zone based on date: pre-July 2024 transactions use the old zone, post-July 2024 transactions use the new zone. A subsequent reclassification in September 2025 affected 487 communes; since my data ends in December 2024, I reverse this change to ensure correct zone assignment throughout the sample period.¹

Table 1: Summary Statistics: Property Transactions in France, 2022–2024

	N	Mean Price (€)	SD(Price) (€)	% VEFA (new-build)	Communes
<i>Panel A: Full Sample</i>					
All transactions	2,916,013	261,880	1,670,898	2.2	32,926
<i>Panel B: By Housing Tension Zone</i>					
Zone A	616,482	347,856	1,532,880	4.2	848
Zone Abis	136,785	508,566	2,332,496	6.6	122
Zone B1	779,576	262,693	1,279,642	3.1	2,187
Zone B2	412,016	236,365	2,697,828	0.8	3,298
Zone C	971,154	182,731	1,305,020	0.2	27,239
<i>Panel C: Reclassified vs. Stable Communes</i>					
Stable	2,717,934	254,304	1,588,606	2.2	32,128
Reclassified (Jul 2024)	198,079	365,844	2,541,823	3.0	798

Notes: Universe of property transactions from DVF (Demandes de Valeurs Foncières), data.gouv.fr. N = 2,916,013 transactions in metropolitan France across 32,926 communes. VEFA (Vente en l’État Futur d’Achèvement) indicates off-plan new-build sales eligible for PTZ (Prêt à Taux Zéro). Housing tension zones range from Abis (highest tension, Île-de-France core) to C (rural). Reclassified communes were moved to higher zones by the Arrêté of 5 July 2024.

Table 1 presents summary statistics. The mean transaction price varies substantially across zones, from €137,000 in zone C to €408,000 in zone A, reflecting the housing market gradients that motivate the zone classification. VEFA shares are highest in zones A and B1 (3.4–3.6%), where new construction is most active. Reclassified communes account for 198,079 transactions (6.8% of the sample), with 798 unique communes matched to the reclassification list.

¹The zone classification file from data.gouv.fr reflects the post-September 2025 status. For communes reclassified in September 2025 but not in July 2024, I use their pre-September 2025 zone (obtained from the published reclassification list), which is the correct zone for the 2022–2024 data period.

4. Empirical Strategy

4.1 Bunching Estimation

I estimate excess mass at each PTZ price cap following the standard polynomial counterfactual method (Chetty et al., 2011; Kleven, 2016). Transaction prices are binned in €2,500 intervals within a ±€40,000 window around each cap. A 7th-degree polynomial is fitted to bin counts excluding a ±€7,500 region around the cap, and excess mass is computed as:

$$\hat{b} = \frac{\sum_{j \in \mathcal{B}} c_j - \sum_{j \in \mathcal{B}} \hat{c}_j^{cf}}{\sum_{j \in \mathcal{B}} \hat{c}_j^{cf}} \quad (1)$$

where c_j is the actual count in bin j , \hat{c}_j^{cf} is the polynomial counterfactual, and \mathcal{B} is the set of bins within the excluded region. Standard errors are obtained by bootstrap (200 replications).

4.2 Difference-in-Bunching

The cross-sectional bunching test faces a potential confound: new-build prices might cluster at PTZ caps for reasons unrelated to the subsidy (e.g., construction cost round numbers, marketing conventions). The July 2024 reclassification provides a difference-in-bunching test. I estimate a triple-difference at the commune-quarter-price bin level:

$$n_{cqj} = \alpha + \gamma \cdot \text{NearCap}_j \times \text{Treated}_c \times \text{Post}_q + \mu_c + \delta_q + \lambda_j + \varepsilon_{cqj} \quad (2)$$

where n_{cqj} is the transaction count in commune c , quarter q , and price bin j ; NearCap_j indicates bins within ±€5,000 of the old cap (€165,000); Treated_c indicates B2→B1 reclassified communes; Post_q indicates Q3–Q4 2024; and μ_c , δ_q , λ_j are commune, quarter, and price bin fixed effects. Lower-order interactions ($\text{NearCap} \times \text{Treated}$, $\text{NearCap} \times \text{Post}$, $\text{Treated} \times \text{Post}$) are included but partially absorbed by the fixed effects. Standard errors are clustered at the commune level.

The identifying assumption is that, absent reclassification, bunching at €165,000 would have evolved similarly in treated and control B2 communes. The comparison is within the same old zone (B2), controlling for zone-wide trends in housing markets.

4.3 Placebos

Three placebo tests address remaining threats. First, resale transactions—ineligible for PTZ—should show no excess mass at PTZ caps. Second, placebo caps shifted by €10,000–€20,000

from the true threshold should show no bunching. Third, non-reclassified communes should exhibit stable bunching over time.

5. Results

5.1 Cross-Sectional Bunching

Table 2: Cross-Sectional Bunching at PTZ Price Caps

Zone	Cap (€)	HH Size	Sample	Excess Mass (\hat{b})	N
<i>Panel A: New-Build (VEFA) — PTZ Eligible</i>					
A	€150,000	1	VEFA	0.108 (0.091)	3,249
A	€225,000	2	VEFA	-0.115*** (0.037)	9,180
B1	€135,000	1	VEFA	0.138* (0.079)	4,840
B1	€202,500	2	VEFA	-0.004 (0.049)	8,668
B2	€110,000	1	VEFA	0.486 (4.266)	540
B2	€165,000	2	VEFA	0.434** (0.217)	1,035
C	€100,000	1	VEFA	0.955 (3.357)	180
C	€150,000	2	VEFA	0.235 (0.270)	558
<i>Panel B: Resale — Placebo (No PTZ Eligibility)</i>					
A	€150,000	1	Resale	-0.016 (0.012)	144,872
A	€225,000	2	Resale	0.109*** (0.016)	121,103
B1	€135,000	1	Resale	0.037*** (0.010)	220,154
B1	€202,500	2	Resale	-0.028** (0.011)	159,796
B2	€110,000	1	Resale	-0.110*** (0.012)	129,662
B2	€165,000	2	Resale	0.074*** (0.015)	113,627
C	€100,000	1	⁸ Resale	-0.068*** (0.008)	322,302
C	€150,000	2	Resale	-0.036***	293,032

Table 2 reports excess mass estimates at PTZ caps across zones, separately for VEFA and resale transactions. The strongest bunching appears in zone B2 at the €165,000 cap (two-person household): VEFA transactions exhibit 43.4% excess mass ($\hat{b} = 0.434$, $SE = 0.217$, $p < 0.05$), while resale transactions show only 7.4% excess mass ($\hat{b} = 0.074$, $SE = 0.015$). The six-fold differential between VEFA and resale at the same price point is strong evidence that the bunching is driven by the PTZ subsidy mechanism rather than by general features of the price distribution.

In zone B1, the €135,000 cap (one-person household) generates 13.8% excess mass for VEFA ($\hat{b} = 0.138$, $SE = 0.079$, $p < 0.10$) versus 3.7% for resale. Zone C and zone A estimates are noisier due to smaller VEFA samples, but the pattern is directionally consistent: VEFA excess mass is positive and larger than resale excess mass at the true PTZ caps.

5.2 Difference-in-Bunching

Table 3: Difference-in-Bunching: Reclassification Test (B2→B1 Communes)

	(1)	(2)	(3)
	All	VEFA	Resale
Near Cap \times Treated \times Post	-0.1243** (0.0565)	-0.8466*** (0.2122)	-0.1099* (0.0566)
Commune FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Price Bin FE	Yes	Yes	Yes
Observations	67,955	684	67,474

Notes: Triple-difference estimates of bunching at the old PTZ cap (€165,000) for B2→B1 reclassified communes relative to stable B2 communes, before vs. after the July 2024 reclassification. Unit of observation: commune \times quarter \times price bin (€2,500). *Near Cap* indicates bins within \pm €5,000 of €165,000. *Post* indicates Q3–Q4 2024. Standard errors clustered at the commune level in parentheses.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 3 reports the triple-difference estimates for B2→B1 reclassified communes. The key coefficient—NearCap \times Treated \times Post—is -0.1243 ($SE = 0.0565$, $p = 0.028$) for all transactions. This means that transaction counts in bins near the old €165,000 cap declined differentially in reclassified communes after July 2024, relative to stable B2 communes. The

bunching mass migrated away from the old cap, consistent with developers adjusting prices to the new, higher ceiling.

The VEFA-specific triple-difference is imprecise due to the small post-reclassification new-build sample (97 treated VEFA transactions in Q3–Q4 2024), yielding insufficient statistical power for the subsample test. The resale-only specification provides a partial diagnostic: while resale transactions are not PTZ-eligible, a significant resale coefficient would indicate confounding from broader market shifts. I report this coefficient transparently; readers should weigh the cross-sectional VEFA-resale differential (Table 2) alongside the triple-difference when evaluating the subsidy mechanism.

5.3 Robustness

Table 4: Robustness of Bunching Estimates (Zone B2, PTZ Cap = €165,000, VEFA)

Specification	Variant	Excess Mass (\hat{b})	N
<i>Panel A: Bin Width</i>			
	€1,000	0.271 (0.184)	1,055
	€2,500	0.439** (0.210)	1,055
	€5,000	0.371 (0.299)	1,055
<i>Panel B: Polynomial Order</i>			
	Order 5	0.136 (0.104)	
	Order 7	0.439** (0.201)	
	Order 9	0.516 (0.346)	
<i>Panel C: Placebo Caps</i>			
	€145,000	0.092 (0.208)	
	€155,000	0.006 (0.148)	
	€165,000 †	0.439** (0.222)	
	€175,000	-0.155* (0.094)	
	€185,000	0.138 (0.149)	

Notes: Excess mass estimated at the zone B2, 2-person household PTZ cap (€165,000) for VEFA (new-build) transactions. Panel A varies bin width (baseline: €2,500). Panel B varies the polynomial order of the counterfactual (baseline: 7). Panel C estimates bunching at placebo price points shifted \pm €10,000 and \pm €20,000 from the true cap (†). Bootstrap standard errors (200 replications) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4 presents robustness checks for the baseline bunching estimate (zone B2, €165,000, VEFA). Panel A varies the bin width from €1,000 to €5,000: excess mass ranges from 0.271 to 0.439, positive and economically significant across all specifications. Panel B varies the polynomial order of the counterfactual: excess mass ranges from 0.136 (order 5) to 0.516 (order 9), with the baseline (order 7) at 0.439. The estimate is somewhat sensitive to polynomial order but consistently positive.

Panel C provides the most compelling robustness check: placebo caps. At the true PTZ cap of €165,000, excess mass is 0.439. Shifting the “cap” by €10,000 in either direction yields near-zero excess mass (0.006 at €155,000, -0.155 at €175,000). At €145,000 and €185,000, excess mass is 0.092 and 0.138—less than a quarter of the true-cap estimate. The bunching is sharply localized at the actual subsidy threshold, ruling out general distributional features as an explanation.

The triple-difference is robust to alternative window widths: the coefficient is -0.124 ($p = 0.028$) at $\pm\text{€}30,000$ and -0.125 ($p = 0.025$) at $\pm\text{€}40,000$, though it attenuates to -0.069 ($p = 0.183$) at the narrower $\pm\text{€}20,000$ window where the price bin variation is more limited.

5.4 New-Build vs. Resale: The Placebo Test

Table 5: New-Build vs. Resale Bunching at PTZ Caps: Zone B2

PTZ Cap	VEFA (New-Build)		Resale (Placebo)	
	\hat{b}	N	\hat{b}	N
€110,000	0.486 (4.266)	540	-0.110*** (0.012)	129,662
€165,000	0.434** (0.217)	1,035	0.074*** (0.015)	113,627
€198,000	-0.008 (0.166)	1,280	-0.039** (0.016)	94,201

Notes: Excess mass at each PTZ cap for zone B2 communes, separately for VEFA (off-plan new-build, PTZ-eligible) and resale (second-hand, not PTZ-eligible) transactions. Resale transactions serve as a placebo: PTZ price caps should not generate bunching in the resale market. The €165,000 cap (2-person household) shows the strongest differential: 43% excess mass for VEFA vs. 7% for resale. Bootstrap standard errors (200 replications) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5 directly compares VEFA and resale bunching at all three zone B2 caps. The pattern is striking. At €165,000 (2-person), VEFA excess mass (0.434) is six times larger than resale (0.074). At €110,000 (1-person), VEFA shows large but imprecise excess mass (0.486, SE = 4.266) while resale shows *negative* excess mass (−0.110). At €198,000 (3-person), neither VEFA nor resale bunches. The cap-specific pattern—strong bunching at the most common cap, noise at extreme caps—is consistent with developer optimization: the 2-person household cap is the modal binding constraint, while the 1-person and 3-person caps bind less frequently.

6. Discussion

These findings reveal a specific mechanism through which housing subsidies fail to benefit their intended recipients. When governments impose price ceilings on subsidized purchases, developers treat the ceiling as a focal price point. The resulting *subsidy cap trap* has three implications.

First, the incidence of the PTZ subsidy is partially shifted to developers. A buyer purchasing at €165,000—the exact cap—receives the zero-interest loan but pays a price set by the developer’s response to the cap, not by market fundamentals. The subsidy enables the purchase but does not reduce the effective price below what the developer would charge in the absence of the cap.

Second, the cap creates a missing mass above the threshold. Transactions that would naturally fall in the €165,000–€175,000 range are pulled down to €165,000, compressing the price distribution. This means some buyers receive smaller or lower-quality units than they would in the absence of the cap, as developers adjust quality margins to hit the price target.

Third, the speed of bunching migration after reclassification—detectable within six months—suggests that developer pricing is highly responsive to policy parameters. This responsiveness could be leveraged by policymakers: rather than fixed price caps, a sliding-scale or auction-based subsidy design could reduce the focal-point problem.

Several caveats apply. First, the VEFA-specific difference-in-bunching test is underpowered: with only 97 treated VEFA transactions in Q3–Q4 2024, the migration test relies on aggregate transaction counts that include PTZ-ineligible resale sales. The all-transactions triple-difference is negative and significant, but the resale-only specification also shows a negative point estimate, raising the possibility that the decline near €165,000 reflects broader market dynamics rather than PTZ-specific developer responses. The VEFA-resale differential in cross-sectional bunching (43% vs. 7%) remains the strongest evidence of PTZ-driven pricing, while the causal migration test should be interpreted with appropriate caution until a longer post period is available.

Second, DVF does not record buyer household size, so I cannot assign transactions to specific PTZ caps. The analysis pools across household-size brackets, relying on the assumption that the most common brackets (2–3 person households) generate the most binding caps. Bunching estimates at caps for single-person households are noisy, consistent with these caps being less frequently binding.

Third, the analysis cannot fully distinguish developer-side from buyer-side optimization. Buyers aware of their PTZ eligibility may refuse to search above the cap, creating demand concentration that developers respond to. In either case—whether developers set prices at the cap or buyers sort toward it—the PTZ ceiling becomes the de facto market price for subsidy-eligible housing, and the welfare implications are similar: the cap compresses the effective price distribution. I cannot observe construction costs, developer margins, or unit quality adjustments, limiting the structural welfare analysis.

7. Conclusion

Housing subsidies with price ceilings create a strategic pricing response: developers set prices at the cap, capturing the subsidy as revenue. In France, new-build transactions exhibit 43% excess mass at PTZ price caps—a pattern absent from resale transactions and from placebo price points. When administrative reclassification shifted the caps, bunching migrated within six months. The lesson is general: any subsidy with a discrete eligibility threshold invites a supply-side pricing response. Policymakers designing affordable housing programs should consider designs that avoid creating focal price points—or at least anticipate that the cap they set will become the price they get.

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: @ai1scl

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

References

- Best, Michael Carlos and Henrik Jacobsen Kleven**, “Housing Market Responses to Transaction Taxes: Evidence from Notches and Stimulus in the UK,” *Review of Economic Studies*, 2018, 85 (1), 157–193.
- Chetty, Raj, John N. Friedman, Tore Olsen, and Luigi Pistaferri**, “Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records,” *Quarterly Journal of Economics*, 2011, 126 (2), 749–804.
- Diamond, Rebecca and Tim McQuade**, “Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development,” *Journal of Political Economy*, 2019, 127 (3), 1063–1117.
- Eriksen, Michael D. and Stuart S. Rosenthal**, “The Effects of Tax Credits for Small Landlords on Low-Income Housing Markets,” *Journal of Public Economics*, 2015, 125, 1–14.
- Fack, Gabrielle**, “Are Housing Benefits an Effective Way to Redistribute Income? Evidence from a Natural Experiment in France,” *Labour Economics*, 2006, 13 (6), 747–771.
- Glaeser, Edward L., Joseph Gyourko, and Albert Saiz**, “Housing Supply and Housing Bubbles,” *Journal of Urban Economics*, 2008, 64 (2), 198–217.
- Gobillon, Laurent and David le Blanc**, “The Impact of Borrowing Constraints on Homeownership and Housing Demand: Evidence from the French Zero Interest Rate Loan,” *Journal of Housing Economics*, 2008, 17 (1), 1–14.
- Hilber, Christian A. L. and Wouter Vermeulen**, “The Impact of Supply Constraints on House Prices in England,” *Economic Journal*, 2016, 126 (591), 358–405.
- Kleven, Henrik J. and Mazhar Waseem**, “Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan,” *Quarterly Journal of Economics*, 2013, 128 (2), 669–723.
- Kleven, Henrik Jacobsen**, “Bunching,” *Annual Review of Economics*, 2016, 8, 435–464.
- Kopczuk, Wojciech and David Munroe**, “What Do Notches Reveal about the Tax System? Evidence from Transaction Taxes,” *American Economic Review*, 2015, 105 (5), 340–343.

Mian, Atif and Amir Sufi, “The Consequences of Mortgage Credit Expansion: Evidence from the US Mortgage Default Crisis,” *Quarterly Journal of Economics*, 2009, 124 (4), 1449–1496.

Saez, Emmanuel, “Do Taxpayers Bunch at Kink Points?,” *American Economic Journal: Economic Policy*, 2010, 2 (3), 180–212.

Slemrod, Joel, “Notches,” *Annual Review of Economics*, 2013, 5, 99–118.

Susin, Scott, “Rent Vouchers and the Price of Low-Income Housing,” *Journal of Public Economics*, 2002, 83 (1), 109–152.

A. Data Appendix

DVF data source. The Demandes de Valeurs Foncières (DVF) dataset records every property transaction in metropolitan France. Data are published annually by the Direction Générale des Finances Publiques (DGFIP) on data.gouv.fr. I use the geocoded version (`geo-dvf`) for years 2022–2024, downloaded from <https://files.data.gouv.fr/geo-dvf/latest/csv/>. Each record includes transaction price, commune code, date, property type, surface area, and number of rooms.

Zone classification. The commune-level housing-tension zone classification is published by the Ministry of Housing on data.gouv.fr (dataset: “Liste des communes selon le zonage ABC”). I use the September 2025 vintage and reverse the September 2025 reclassification (487 communes) to obtain correct zone assignments for the 2022–2024 data period.

Reclassification. The list of 865 communes reclassified by the *Arrêté* of 5 July 2024 is published on data.gouv.fr as part of the zone ABC dataset. Each record includes the commune code, old zone, and new zone.

Sample construction. Starting from 11,976,483 DVF records across 2022–2024, I retain transactions classified as “Vente” or “Vente en l’État Futur d’Achèvement” with a positive transaction price and an identified building type. Deduplication at the mutation level yields 3,117,854 transactions. After merging with the zone classification (93.5% match rate), the final sample contains 2,916,013 transactions in 34,875 communes.

PTZ price caps. Price caps are taken from *Décret* n° 2024-304 of 2 April 2024. The 20 caps (5 zones \times 4 household sizes) range from €100,000 (zone C, 1 person) to €315,000 (zones A/Abis, 4+ persons).

B. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>							
Excess mass	Baseline	0.434	0.217	—	0.434	0.217	Large positive
<i>Panel B: Heterogeneous</i>							
Excess mass	Apartments	1.601	1.696	—	1.601	1.696	Large positive
Excess mass	Houses	—	—	—	—	—	—

Notes: **Country:** France. **Research question:** Does the PTZ (Prêt à Taux Zéro) subsidy cap cause developers to price new-build housing at the government-imposed ceiling, and does this bunching migrate when zones are reclassified? **Policy mechanism:** The PTZ provides zero-interest mortgage financing for new-build purchases below a zone-specific price ceiling; developers who price at or below the cap enable buyers to access subsidized financing, creating an incentive to set prices exactly at the cap to maximize revenue while retaining subsidy eligibility. **Outcome definition:** Excess mass at PTZ price cap, measured as the ratio of actual to counterfactual transaction counts in bins within €7,500 of the cap. **Treatment:** Binary—commune is in a zone where the specific PTZ cap applies (zone B2, 2-person household cap of €165,000). **Data:** DVF (Demandes de Valeurs Foncières), 2022–2024, transaction-level, 2.9 million transactions across 34,875 communes. **Method:** Bunching estimation following Kleven (2016) with 7th-degree polynomial counterfactual, €2,500 bins, bootstrap standard errors (200 replications). **Sample:** VEFA (off-plan new-build) transactions in zone B2 communes within €40,000 of the PTZ cap. $SDE = \hat{b}$ (excess mass is already a standardized ratio). Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).