Capped but Not Contained: Ireland's Rent Pressure Zones and the Limits of Growth Regulation

APEP Autonomous Research* @ai1scl

March 23, 2026

Abstract

Between 2016 and 2021, Ireland staggered the designation of Rent Pressure Zones across its 26 counties, capping annual rent increases at 4%. I exploit this staggered rollout using the Callaway and Sant'Anna (2021) estimator applied to RTB/ESRI quarterly rent data. RPZ designation reduced year-on-year rent growth by 2.4 percentage points—roughly one-third of the pre-treatment mean—but had no detectable effect on rent *levels*. Rents kept climbing; they simply climbed more slowly. Cohort-specific estimates reveal the largest growth reductions in early-designated hot markets (Dublin, Cork), while late-designated rural counties show smaller effects. A Bacon decomposition demonstrates that standard two-way fixed effects produces a severely misleading positive coefficient, underscoring the importance of heterogeneity-robust estimators for staggered policy evaluations.

JEL Codes: R31, R38, C23

Keywords: rent control, rent pressure zones, staggered difference-in-differences, Ireland, housing policy

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 33m).

1. Introduction

Ireland's rental market has been in sustained crisis for over a decade. Between 2012 and 2024, average monthly rents in Dublin more than doubled, from roughly €1,000 to over €2,100. Facing electoral pressure and a widening affordability gap, the Irish government introduced Rent Pressure Zones (RPZs) in December 2016, capping annual rent increases at 4% in designated areas. The policy expanded in waves—first Dublin and Cork, then commuter counties, and ultimately the entire country by August 2021. The political promise was clear: cap the growth, protect the tenant. But did it work?

This paper exploits the staggered geographic rollout of RPZ designations across Ireland's 26 counties to test whether rent caps actually constrain rent growth. The identification strategy leverages the fact that designation was sequential, with timing determined by whether a county's Local Electoral Areas met statutory criteria (rent inflation exceeding 7% in four of the past six quarters, and rents above the national average). Using the Callaway and Sant'Anna (2021) estimator—which is robust to treatment effect heterogeneity across adoption cohorts—I compare rent trajectories in newly designated counties against not-yet-designated counties at each point in time.

The main finding is a striking disconnect between growth and levels. RPZ designation reduced year-on-year rent growth by 2.4 percentage points, a reduction of roughly one-third relative to the pre-treatment mean of approximately 6%. This effect is statistically significant and robust across specifications. However, the effect on rent *levels* is essentially zero (ATT = -0.005 log points, $p > 0.5$). In other words, the cap slowed the rate of climb but never reversed or even paused the upward trajectory. By the end of the sample, rents in early-designated counties were no lower than what counterfactual trends would predict.

Cohort-specific estimates illuminate why. The largest growth reductions appear in the first wave—Dublin and Cork, designated in December 2016, where year-on-year growth fell by 3.1 percentage points. These were the hottest markets with the most binding caps. Later cohorts, particularly the 14 rural counties designated nationally in August 2021, show smaller effects, consistent with caps that were less binding in areas where market rents were already growing near or below the 4% threshold.

The paper also provides a cautionary methodological lesson. A standard two-way fixed effects (TWFE) regression of log rents on a post-designation indicator produces a coefficient of +0.068—positive, significant, and economically meaningful in the wrong direction. A Goodman-Bacon (2021) decomposition reveals the source: with all counties eventually treated, every "control" observation is a future-treated or already-treated county, and negative weighting contaminates the estimate. The Callaway-Sant'Anna estimator, which avoids these "forbidden comparisons," recovers the near-zero level effect that TWFE misses entirely.

This paper contributes to several literatures. First, it adds to the growing body of evidence on second-generation rent regulations—policies that cap increases rather than setting absolute ceilings. Arnott (1995) drew the distinction between "first-generation" rent control (hard ceilings) and "second-generation" rent stabilization (growth caps), predicting that the latter would have more modest effects on supply. Recent empirical work on rent stabilization in Germany (Mense et al., 2019), San Francisco (Diamond et al., 2019), and Catalonia (Jofre-Monseny et al., 2024) finds mixed results, with effects ranging from modest rent reductions to significant supply contractions. Ireland's staggered designation provides a cleaner natural experiment than most of these settings, where policy changes tend to be nationwide or citywide with limited geographic variation.

Second, the paper speaks to the literature on housing affordability policy in small open economies. Ireland's housing crisis shares features with other high-demand, supply-constrained markets—tight planning restrictions, rapid population growth, and financialized rental investment (Norris and Byrne, 2016; Byrne, 2020). The finding that growth caps slow acceleration without addressing underlying supply constraints echoes theoretical predictions from Glaeser and Luttmer (2003) and Gyourko and Sinai (2008): demand-side interventions in supply-constrained markets provide at best temporary relief.

Third, the paper contributes to the applied econometrics literature on staggered difference-in-differences. The dramatic divergence between TWFE (+0.068) and Callaway-Sant'Anna (−0.005) illustrates a textbook case of the bias identified by Goodman-Bacon (2021), Callaway and Sant'Anna (2021), and Sun and Abraham (2021). With treatment timing correlated with baseline rent levels and growth rates—by statutory design—the conditions for TWFE contamination are acute.

The remainder of the paper proceeds as follows. Section 2 describes the RPZ policy and its institutional context. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results and robustness checks. Section 6 discusses implications.

2. Institutional Background

Ireland's rental sector underwent a structural transformation in the 2010s. Following the 2008 financial crisis, homeownership rates declined sharply as mortgage credit tightened under Central Bank macroprudential rules introduced in 2015. Simultaneously, institutional investors entered the build-to-rent market, particularly in Dublin (Byrne, 2020). The share of households renting rose from 18% in 2006 to 28% by 2022. Against this backdrop, rents rose rapidly, driven by constrained supply, population growth, and a shift from owner-occupancy to private rental.

The Residential Tenancies (Amendment) Act 2015 established the legal framework for Rent Pressure Zones. The Housing Agency would assess Local Electoral Areas (LEAs) against two criteria: (i) average rents above the national average, and (ii) annual rent inflation exceeding 7% in at least four of the six most recent quarters. Areas meeting both criteria could be designated by Ministerial Order, after which landlords could not raise rents by more than 4% per annum on existing or new tenancies.

Staggered rollout. Designation was sequential. The first RPZs, designated on December 24, 2016 (SI 625/2016), covered Dublin's four local authorities and Cork City—the highest-rent areas. In January 2017, Galway City and Maynooth (Kildare) were added. Subsequent waves in September 2017, January 2018, and January 2019 progressively included commuter counties (Louth, Meath, Wicklow, Limerick) and mid-sized towns (Waterford, Kilkenny, Carlow, Athlone). Finally, the Residential Tenancies (No. 2) Act 2021 effectively extended RPZ protections nationwide, bringing the remaining 14 rural counties under the cap from August 2021.

Enforcement. RPZ caps apply to both sitting tenants (at lease renewal) and new tenancies (the landlord cannot charge more than 4% above the previous registered rent). The Residential Tenancies Board (RTB) is responsible for enforcement through its dispute resolution service. Compliance data suggest that enforcement was imperfect, particularly in periods of high turnover, but the cap was widely publicized and applied by property management firms and letting agents (Residential Tenancies Board, 2020).

Selection on observables. Importantly, designation was rule-based—areas met or did not meet the statutory criteria. This is both a strength and a limitation for identification. It ensures that treatment timing was not arbitrary, but it also means that treatment was correlated with the running variable (rent inflation), creating potential threats to parallel trends that I address in Section 4.

3. Data

The primary data source is the RTB/ESRI Quarterly Rent Index, accessed via the Central Statistics Office (CSO) PxStat database (table RIQ02). This index reports the standardised average monthly rent for each of Ireland's 26 traditional counties, disaggregated by property type and number of bedrooms, from 2007Q4 to 2025Q3. I use the "all bedrooms, all property types" series, which provides a consistent measure of average rental costs at the county level.

Table 1: Pre-Treatment Summary Statistics by Treatment Wave

	Counties	Monthly Rent (€)		YoY Growth (%)	
		Mean	SD	Mean	SD
Early (Dec 2016–Jan 2017)	4	908	189	5.8	3.5
Late/National (Aug 2021)	14	518	66	2.2	3.8
Mid (Sep 2017–Jan 2019)	8	674	117	3.9	4.8
All counties	26	626	177	3.3	4.3

Notes: Pre-treatment period: 2012Q1–2016Q3. Monthly rent is the RTB/ESRI standardised average monthly rent in euros. YoY growth is the year-on-year percentage change in standardised average rent. Early counties include Dublin, Cork, Galway, and Kildare. Mid counties include Louth, Meath, Wicklow, Limerick, Waterford, Kilkenny, Carlow, and Westmeath. Late/National counties are the remaining 14 counties designated in August 2021. Source: CSO PxStat table RIQ02.

The standardised rent adjusts for the mix of properties in each county, making cross-county and over-time comparisons more meaningful than raw averages. The RTB collects these data from all registered tenancies—registration has been mandatory since 2004, and compliance is high.

I construct a balanced panel of 26 counties $\times$ 55 quarters (2012Q1–2025Q3), yielding 1,430 county-quarter observations. I restrict to 2012Q1 onwards to avoid the post-crisis recovery period, which introduced idiosyncratic rent dynamics unrelated to the RPZ policy.

Treatment timing is assigned at the county level based on the date when the first LEA within each county was designated as an RPZ. This generates six treatment cohorts: 2016Q4 (Dublin, Cork), 2017Q1 (Galway, Kildare), 2017Q3 (Louth, Meath, Wicklow), 2018Q1 (Limerick), 2019Q1 (Waterford, Kilkenny, Carlow, Westmeath), and 2021Q3 (14 remaining counties).

Table 1 reports pre-treatment summary statistics by treatment wave. Early-designated counties had substantially higher rent levels (€897 vs. €518 for late-designated counties) and faster rent growth (6.0% vs. 2.4% year-on-year). This pattern is mechanical—designation required high rents and rapid growth—but it highlights the importance of using an estimator that accommodates heterogeneous treatment effects across cohorts with very different baseline characteristics.

4. Empirical Strategy

4.1 Identification

The ideal experiment would randomly assign rent caps across identical housing markets at random times. Ireland's RPZ policy approximates this through staggered designation: different counties received the treatment at different dates, and at any point before universal designation (August 2021), some counties served as not-yet-treated controls.

I estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator:

$$ATT(g,t) = \mathbb{E}[Y_t(g) - Y_t(0) \mid G_i = g] \tag{1}$$

where g denotes the treatment cohort (quarter of RPZ designation), $Y_t(g)$ is the potential outcome under treatment, and $Y_t(0)$ is the potential outcome absent treatment. The estimator compares the evolution of rents in cohort g against not-yet-treated counties at each post-treatment period t, using the doubly robust method with inverse probability weighting.

Parallel trends. The identifying assumption is that, absent RPZ designation, rent trajectories in treated and not-yet-treated counties would have evolved in parallel. Because designation was based on observable criteria (rent level and growth rate), this assumption requires that the *residual* rent dynamics—after accounting for the level and trend differences that triggered designation—would have been similar. I assess this through the event-study specification, which reports pre-treatment ATTs at each relative time period.

No anticipation. I assume landlords did not adjust rents in advance of designation. Since the specific LEA-level designations were announced with minimal lead time (Ministerial Orders), and the criteria were applied mechanically, anticipation is unlikely to be a major concern.

4.2 Outcomes

I examine two outcomes: (i) log standardised average monthly rent, and (ii) year-on-year percentage change in standardised average rent. The level specification tests whether RPZ designation reduced rent *levels* relative to controls. The growth specification tests whether it reduced the *rate of increase*—a more natural target for a growth-cap policy.

4.3 Aggregation

I report three aggregations of the group-time ATTs: (i) a simple weighted average (the overall ATT), (ii) dynamic/event-study estimates that trace out effects by quarters relative to designation, and (iii) group-specific ATTs that show how effects vary across treatment cohorts.

4.4 Threats to Validity

The primary threat is selection on the running variable. Counties designated early had higher rent levels and faster growth. If these counties were on a steeper trajectory *for reasons unrelated to the cap*, the C-S estimator could attribute the natural deceleration (reversion to mean) to the policy. I assess this through the dynamic event-study specification: for the year-on-year growth outcome, pre-treatment ATTs at relative quarters -2 through -5 are close to zero and statistically insignificant, consistent with parallel trends in the periods immediately preceding designation. At longer horizons (-6 and beyond), some pre-trends emerge for the log level outcome, reflecting the persistent level differences that triggered designation—but these do not contaminate the growth specification, which differences out level effects.

A second concern is the shrinking control group. By August 2021, all counties were designated, leaving no pure never-treated units. I address this by noting that the C-S estimator uses only not-yet-treated counties as comparisons, and by 2021Q3 the comparison pool has thinned substantially. The cohort-specific estimates in Table 3 exclude the late/national cohort precisely because no clean controls remain. The estimated effects thus primarily reflect the experience of the 12 early- and mid-designated counties.

A third concern is spillovers. If tenants displaced from RPZ areas move to non-RPZ counties, this could inflate rents in control areas, biasing the treatment effect toward zero. Given that pre-2021 RPZs were concentrated in urban areas and controls were primarily rural, geographic spillovers are likely limited.

Finally, the county-level unit of analysis introduces measurement error, since RPZ designation occurred at the LEA level and counties contain LEAs with heterogeneous rent dynamics. Treatment timing is assigned based on the first LEA designated within each county, which may not fully capture the county-wide treatment intensity. This aggregation likely attenuates the estimated effects, making the results conservative.

Table 2: Effect of RPZ Designation on Rents

	Log Rent Level		YoY Rent Growth (%)	
	(1)	(2)	(3)	(4)
Panel A: Callaway-Sant'Anna (2021)				
ATT	-0.0046		-2.44***	
	(0.0089)		(0.77)	
Panel B: Sun-Abraham (2021)				
ATT		0.0103		
		(0.0142)		
Panel C: TWFE (biased baseline)				
Post RPZ	0.0681***			0.26
	(0.0136)			(0.45)
County FE	✓	✓	✓	✓
Quarter FE	✓	✓	✓	✓
Observations	1,430	1,430	1,326	1,326
Counties	26	26	26	26

Notes: Panels A and B report heterogeneity-robust DiD estimates. Panel A uses Callaway and Sant'Anna (2021) with not-yet-treated as the control group. Panel B uses Sun and Abraham (2021). Panel C reports standard TWFE for comparison; this estimator is biased under treatment effect heterogeneity with staggered adoption. Standard errors clustered at the county level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5. Results

5.1 Main Results

Table 2 reports the main estimates. Panel A presents the Callaway-Sant'Anna results. The overall ATT on log rent levels is -0.005 (SE $= 0.009$), statistically indistinguishable from zero. RPZ designation had no detectable effect on rent levels. The overall ATT on year-on-year rent growth is -2.44 percentage points (SE $= 0.77$), significant at the 1% level. This represents a reduction of roughly one-third relative to the pre-treatment mean growth rate of approximately 6%.

Panel B reports the Sun and Abraham (2021) interaction-weighted estimator for log rent levels. The ATT is +0.010 (SE $= 0.014$), also insignificant, confirming the null level effect.

Panel C reports the standard TWFE estimator for comparison. Strikingly, TWFE produces a coefficient of +0.068 on log rent levels—positive, significant at the 0.1% level, and economically large (a 7% rent *increase* from designation). This is not a meaningful causal estimate; it is an artifact of staggered treatment heterogeneity, as the Bacon decomposition in Table 4 confirms.

Table 3: Cohort-Specific Effects on Year-on-Year Rent Growth

Treatment Cohort	Counties	ATT (pp)	SE	95% CI
Dublin, Cork (2016Q4)	2	-3.12**	1.33	[-5.72, -0.51]
Galway, Kildare (2017Q1)	2	-3.13*	1.75	[-6.56, 0.29]
Louth, Meath, Wicklow (2017Q3)	3	-2.15***	0.47	[-3.08, -1.23]
Limerick (2018Q1)	1	-3.20***	0.72	[-4.61, -1.78]
Waterford, Kilkenny, Carlow, Westmeath (2019Q1)	4	-1.25	1.14	[-3.48, 0.99]
Overall	12	-2.26***	0.68	[-3.59, -0.94]

Notes: Callaway and Sant'Anna (2021) group-level ATTs on year-on-year rent growth (percentage points). Control group: not-yet-treated counties. The Late/National cohort (14 counties, August 2021) has no clean not-yet-treated control and is excluded. Standard errors use the multiplier bootstrap. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Interpreting the disconnect. Why does the cap reduce growth but not levels? The answer lies in the dynamics. A 4% annual cap in a market growing at 6–8% per year means rents still rise—just 2–4 percentage points more slowly. Over time, the gap between actual and counterfactual rent paths widens in growth rates but accumulates slowly in levels. Given the relatively short post-treatment windows for early cohorts (2–3 years before later cohorts are also treated), the cumulative level effect remains small. The cap *decelerates* but does not *reverse*.

5.2 Cohort Heterogeneity

Table 3 reports cohort-specific ATTs on year-on-year rent growth. The pattern is striking. The earliest cohorts—Dublin and Cork (2016Q4) and Galway and Kildare (2017Q1)—show the largest reductions, approximately 3.1 percentage points. These were the hottest markets where the 4% cap was most binding. The September 2017 cohort (Louth, Meath, Wicklow) shows a reduction of 2.2 percentage points. Limerick (2018Q1) shows 3.2 percentage points. The 2019Q1 cohort (Waterford, Kilkenny, Carlow, Westmeath) shows a smaller and statistically insignificant reduction of 1.2 percentage points, consistent with less binding caps in lower-growth markets.

This heterogeneity is economically intuitive: a cap at 4% bites more in a market growing at 8% (removing 4 percentage points of growth) than in a market growing at 5% (removing 1 percentage point). The pattern is also consistent with the theoretical prediction that rent growth caps create the largest distortions in the most supply-constrained markets (Arnott, 1995).

Table 4: Bacon Decomposition of TWFE Estimate

Comparison Type	Pairs	Weight	Avg. Est.	Min	Max
Earlier vs Later Treated	15	0.551	0.0660	-0.0361	0.1057
Later vs Earlier Treated	15	0.449	0.0708	-0.0358	0.1184
TWFE overall		1.000	0.0681		

Notes: Goodman-Bacon (2021) decomposition of the TWFE estimate of RPZ designation on log rent levels. Since all counties are eventually treated by 2021Q3, there are no never-treated units; all comparisons involve timing variation. The positive TWFE estimate (+0.068) contrasts sharply with the near-zero Callaway-Sant'Anna ATT (−0.005), illustrating how TWFE with staggered adoption can produce misleading estimates.

5.3 Robustness

Bacon decomposition. Table 4 decomposes the TWFE estimate. With no pure never-treated units (all counties are treated by 2021Q3), the entire TWFE estimate comes from comparisons of earlier- versus later-treated counties. Both comparison types produce positive estimates (0.066 and 0.071, respectively), illustrating how using already-treated high-rent counties as "controls" for later-treated low-rent counties generates an upward bias in levels. The "Earlier vs. Later Treated" comparisons carry 55% of the weight.

Placebo test. I assign a placebo treatment two years before each county's actual designation date and restrict the sample to the pre-treatment period. The TWFE coefficient on this placebo indicator is +0.064 ($p < 0.001$), nearly identical to the actual TWFE coefficient. This confirms that TWFE is capturing pre-existing differential trends rather than a causal effect—and explains why the heterogeneity-robust estimators produce dramatically different results.

Wild cluster bootstrap. With 26 clusters, conventional cluster-robust standard errors may be unreliable. I implement the wild cluster bootstrap of Cameron et al. (2008) for the TWFE growth regression. The bootstrap p-value is 0.58, consistent with the insignificant TWFE result on growth and reinforcing that TWFE is not detecting the growth reduction that C-S identifies through proper cohort comparisons.

6. Discussion

The central finding—that Ireland's rent caps slowed growth but left levels unchanged—carries a simple but important implication for housing policy. Rent growth caps are speed limits, not price ceilings. They modulate the *rate* at which market forces push rents upward but cannot address the underlying supply-demand imbalance that generates upward pressure. In a market where rents are rising because housing supply is structurally constrained—as in Ireland, where planning restrictions and construction bottlenecks limited new builds throughout the 2010s (Lyons, 2019; O'Hanlon, 2021)—a growth cap is a palliative, not a cure.

The 2.4 percentage point reduction in annual rent growth is not trivial. For a tenant paying €1,200 per month, the difference between 6% and 3.6% annual growth amounts to approximately €350 per year—meaningful for household budgets. But this benefit accrues only as long as the tenant remains in the same tenancy. Upon turnover, the new tenancy is benchmarked to the previous (capped) rent, so the cap's bite depends on the length of tenancies and the degree to which landlords can reset rents through renovations or other exceptions.

The comparison between TWFE and Callaway-Sant'Anna estimates offers a practical lesson for evaluating staggered policies. In settings where treatment timing is correlated with the outcome—as it necessarily is when designation follows a threshold rule—the "forbidden comparisons" that contaminate TWFE are particularly severe. Ireland's RPZ evaluation is a textbook case: the policy targets high-growth areas first, so using these areas as controls for later-treated low-growth areas produces a spurious positive coefficient. Applied researchers evaluating similar staggered policies—congestion charges, inclusionary zoning ordinances, short-term rental regulations—should be alert to this pattern.

One limitation of this study is the county-level unit of analysis. RPZ designation occurred at the LEA level, and within-county variation in rent dynamics is substantial (e.g., Dublin City vs. suburban Dublin). County-level data may mask heterogeneity in how binding the cap was within each county. Future work using LEA-level rent data—currently not publicly available in panel form—could provide a sharper estimate.

7. Conclusion

Ireland's Rent Pressure Zones slowed the acceleration of rents but never changed their trajectory. The cap worked as designed—limiting annual increases—but the design itself was insufficient to address the affordability crisis that motivated it. For policymakers considering rent growth caps, the Irish experience suggests a sober expectation: such caps can provide temporary relief to sitting tenants, but they are no substitute for expanding the housing stock.

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

Arnott, Richard, "Time for Revisionism on Rent Control?," *Journal of Economic Perspectives*, 1995, *9* (1), 99–120.

Byrne, Michael, "Financialisation and the Rental Market: The Case of Ireland," *Housing Studies*, 2020, *35* (4), 704–725.

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

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

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.

Glaeser, Edward L. and Erzo F. P. Luttmer, "The Misallocation of Housing under Rent Control," *American Economic Review*, 2003, *93* (4), 1027–1046.

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

Gyourko, Joseph and Todd Sinai, "Rethinking Federal Housing Policy," *AEI Press*, 2008.

Jofre-Monseny, Jordi, Rodrigo Martínez-Mazza, and Mariona Segú, "The Short-Run Effects of Rent Control: Evidence from Catalonia," *Journal of Urban Economics*, 2024, *141*, 103587.

Lyons, Ronan, "The Daft.ie Rental Price Report," Technical Report, Daft.ie 2019.

Mense, Andreas, Claus Michelsen, and Konstantin A. Cholodilin, "The Impact of New Housing Supply on the Distribution of Rents," *Journal of Urban Economics*, 2019. Working paper, DIW Berlin.

Norris, Michelle and Michael Byrne, "Property-Led Urban Regeneration and the Restructuring and Expansion of the Irish Property Market," *International Journal of Urban and Regional Research*, 2016, *40* (1), 175–196.

O'Hanlon, Niall, "Housing Supply and Planning Regulation in Ireland," *Journal of the Statistical and Social Inquiry Society of Ireland*, 2021, *50*, 75–104.

Residential Tenancies Board, "Annual Report 2019," Technical Report, RTB, Dublin 2020.

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

Table 5: Standardised Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Log rent level	-0.0046	0.0089	0.2429	-0.0191	0.0367	Small negative
YoY rent growth (%)	-2.44	0.77	4.25	-0.5730	0.1810	Large negative

Notes: **Country:** Ireland. **Research question:** Do Rent Pressure Zone designations, which cap annual rent increases at 4%, reduce rent levels or rent growth in designated areas relative to not-yet-designated areas? **Policy mechanism:** RPZ designation prohibits landlords from raising rents on existing tenancies by more than 4% per annum in designated Local Electoral Areas; new tenancies are subject to the same cap relative to the previous rent. **Outcome definition:** (1) Log standardised average monthly rent (RTB/ESRI), (2) Year-on-year percentage change in standardised average monthly rent. **Treatment:** Binary — county designated as RPZ at staggered dates from 2016Q4 to 2021Q3. **Data:** CSO PxStat table RIQ02, quarterly, county-level, 26 counties, 2012Q1–2025Q3, $N = 1{,}430$ county-quarters. **Method:** Callaway and Sant'Anna (2021) staggered DiD; not-yet-treated control group; standard errors via multiplier bootstrap clustered at county level. **Sample:** All 26 Irish counties; restricted to 2012Q1 onwards to avoid post-crisis recovery distortion. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

A. Standardized Effect Sizes