Capped but Not Crowded Out: The Benefit Cap, Housing Markets, and the Parallel Trends Trap

APEP Autonomous Research* @SocialCatalystLab

March 18, 2026

Abstract

England's November 2016 benefit cap reduction—from £26,000 to £20,000/£23,000—tripled the capped population from 20,000 to 68,000 households. Policymakers and advocates argued this would drive a homelessness crisis. Exploiting the enormous cross-local-authority variation in cap exposure (from 0.09 to 3.93 capped households per 1,000 population), I estimate continuous-treatment difference-in-differences models on 278 English local authorities observed annually from 2013 to 2018. The point estimate is positive but statistically insignificant, and an event study reveals significant pre-trends: high-intensity authorities were already experiencing faster temporary accommodation growth three years before the reform. The parallel trends assumption fails, and the naïve coefficient likely captures housing market pressure—not cap-induced displacement. This result illustrates how confounding policy targeting with outcome trends can generate misleading attributions of welfare reform effects.

JEL Codes: H53, I38, R21
Keywords: benefit cap, homelessness, temporary accommodation, parallel trends, housing policy, welfare reform

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 10h 1m).

1. Introduction

On November 7, 2016, the UK government lowered the benefit cap from £26,000 to £20,000 per year outside London and £23,000 in London, tripling the number of households whose total welfare payments were capped. Birmingham alone saw 3,079 households hit by the new ceiling—more than many entire counties. The stated purpose was fiscal: reduce government spending and sharpen work incentives. But critics argued the real consequence would be housing displacement, as capped households could no longer cover their rent and would be forced into homelessness. If true, the Department for Work and Pensions would be saving on benefits while local authorities paid for expensive temporary accommodation—a fiscal shell game rather than genuine savings.

This paper tests that claim directly. Using the enormous variation in cap exposure across 278 English local authorities—ranging from 0.09 to 3.93 capped households per 1,000 population—I estimate whether authorities with greater exposure to the benefit cap reduction experienced larger increases in temporary accommodation placements. The continuous-treatment difference-in-differences design exploits the fact that cap intensity is predetermined: it depends on the pre-existing geographic distribution of high-benefit claimants, driven by local housing costs and household composition rather than by the reform itself.

The naïve estimate suggests a positive effect, but the identification does not survive scrutiny. The interaction of cap intensity with a post-reform indicator yields a positive coefficient of 0.312 temporary accommodation households per 1,000 population (standard error 0.177), consistent with the displacement hypothesis. However, this estimate is not credibly causal. An event study reveals significant pre-treatment coefficients: relative to 2016, the cap intensity interaction is −0.554 in 2012 ($p < 0.001$), −0.416 in 2013 ($p = 0.004$), −0.270 in 2014 ($p = 0.017$), and −0.118 in 2015 ($p = 0.111$). The post-treatment coefficients—0.030 in 2017 and 0.029 in 2018—show no break in the pre-existing trend. High-cap-intensity authorities were already experiencing faster temporary accommodation growth well before the cap was reduced.

This pre-trend is robust across specifications. Placebo tests that assign treatment at 2014 or 2015 yield larger and more significant coefficients than the actual reform date. Excluding London—which has a separate, higher cap threshold—eliminates the positive effect entirely ($\hat{\beta} = -0.047$, $p = 0.846$). Adding region-by-year fixed effects to absorb differential regional trends also drives the estimate to zero ($\hat{\beta} = -0.065$, $p = 0.753$). The pattern is consistent: the cross-sectional correlation between cap bite and temporary accommodation growth reflects common underlying causes—tightening housing markets, declining social housing stock, and rising rents—rather than a causal chain from the benefit cap to homelessness.

This finding contributes to the literature on welfare reform and housing in three ways. First, it provides a methodological caution for evaluating geographically-targeted welfare policies. The benefit cap was disproportionately applied in areas with high housing costs (Beatty and Fothergill, 2016), which are precisely the areas where homelessness pressures were already mounting (Fitzpatrick et al., 2019). Any design that uses cross-area variation in policy exposure must contend with this selection—the same housing market conditions that generate high cap exposure also drive rising homelessness independently.

Second, the result speaks to the broader debate about the fiscal incidence of welfare reform. Reeves and Loopstra (2017) found the benefit cap increased employment among affected households but raised food bank usage. Brewer et al. (2019) documented that the transition to Universal Credit created housing payment shortfalls. The present findings suggest that while these individual-level effects may be real, the aggregate-level consequence for local authority temporary accommodation burdens is confounded with pre-existing trends that any cross-sectional design would attribute to the policy.

Third, the paper illustrates a general identification problem in continuous-treatment designs: when policy intensity correlates with outcome trends for structural reasons. The benefit cap is a canonical example where exposure varies geographically and affects a readily measured downstream outcome. The failure of parallel trends here—with 278 units, seven years, and enormous variation in treatment intensity—raises caution about similar continuous-treatment designs applied to other place-based welfare reforms (Bronchetti et al., 2023; Clarke and Martorell, 2020). This does not mean the question is unanswerable, only that simple aggregate cross-area designs are vulnerable to this particular confound. Individual-level matched designs, as in the DWP's own evaluation (Department for Work and Pensions, 2019), avoid the ecological correlation by conditioning on household characteristics rather than geographic exposure.

The rest of the paper proceeds as follows. Section 2 describes the institutional setting and the benefit cap reduction. Section 3 presents the data sources and panel construction. Section 4 details the empirical strategy and identification assumptions. Section 5 presents results, including the event study and robustness checks. Section 6 discusses implications. Section 7 concludes.

2. Institutional Background

The benefit cap. The UK benefit cap was introduced in April 2013 as part of the Welfare Reform Act 2012. It limits total welfare payments—including Housing Benefit, Income Support, Jobseeker's Allowance, and Child Benefit—to a household-level maximum. Households with at least one member in employment above a minimum-hours threshold are exempt, as are those receiving disability benefits, War Widows' Pension, or Working Tax Credit. The original cap was set at £26,000 per year (£500 per week), roughly the median earned income.

The 2016 reduction. The Welfare Reform and Work Act 2016 authorized a substantial reduction. From November 7, 2016, the cap fell to £20,000 per year (£384.62 per week) outside London and £23,000 (£442.31 per week) in London, a 23% and 12% cut respectively. The rollout was implemented Jobcentre by Jobcentre between November 2016 and January 2017.

The reduction dramatically expanded the capped population. Under the original £26,000 cap, approximately 20,000 households were affected at any point in time. After the reduction, this rose to approximately 68,000 by May 2017—a more-than-threefold increase. The geographic distribution was highly uneven. Birmingham had 3,079 capped households; Brent had 1,319; Hackney had 1,040. At the other extreme, the Isle of Wight had just 13 capped households, Calderdale 22, and Harborough 17.

Temporary accommodation. When a household becomes homeless and presents to a local authority, the authority has a duty under Part 7 of the Housing Act 1996 to provide temporary accommodation (TA) if the household is in priority need. TA is expensive—the National Audit Office estimated that the average cost to local authorities was £18,000 per household per year in 2017, rising to over £25,000 in London (National Audit Office, 2017). The combination of declining social housing stock and rising private rents meant that temporary accommodation use was already climbing before the cap reduction: from 50,430 households in March 2012 to 71,670 in March 2016, an increase of 42%.

The displacement hypothesis. If the cap forces households to leave their current accommodation—because Housing Benefit no longer covers the rent—some will fail to find alternative housing and present as homeless, generating additional demand for temporary accommodation. The displacement effect should be proportional to the number of newly capped households in each local authority. This cross-authority variation in cap bite forms the basis of the identification strategy.

3. Data

The analysis combines three data sources. The primary outcome is from the MHCLG statutory homelessness returns (Table 784), which records the number of households in temporary accommodation at the end of each financial year (31 March), disaggregated

Table 1: Summary Statistics, Pre-Treatment (2012–2016)

	Mean	SD	p10	Median	p90
TA households per 1,000 pop	0.98	2.20	0.04	0.24	2.55
Capped HH per 1,000 pop	1.00	0.61	0.42	0.86	1.72
Claimant rate per 1,000	13.22	8.11	4.72	11.34	24.33
Population (thousands)	178.69	119.24	84.74	138.69	317.88
Local authorities			278		
LA $\times$ year observations			1367		

Notes: Pre-treatment period is financial years ending March 2012–2016. TA households per 1,000 population is from MHCLG Table 784. Capped households per 1,000 population measured at May 2017 (first post-reduction snapshot). Claimant rate is the annual average JSA/UC claimant count per 1,000 population from NOMIS. Sample restricted to 278 English local authorities with non-missing data in all periods.

by local authority, covering 2009/10 through 2017/18 (Ministry of Housing, Communities and Local Government, 2018). These returns were collected under the P1E form from all English local housing authorities until the transition to H-CLIC in April 2018. I restrict the panel to 2012–2018 (financial years ending March 2012 through March 2018), providing five pre-treatment and two post-treatment years.

Treatment intensity comes from the DWP Benefit Cap Statistics for May 2017, the first comprehensive post-reduction quarterly snapshot (Department for Work and Pensions, 2017). This reports the point-in-time number of capped households by local authority. I normalize by mid-year population from NOMIS to create a cap intensity measure: capped households per 1,000 residents.

Control variables include the annual average JSA/UC claimant count per 1,000 population from NOMIS, which proxies for local labor market conditions.

3.1 Summary Statistics

The final panel contains 278 English local authorities observed over seven years (1,870 authority-year observations). The pre-treatment mean of the temporary accommodation rate is 0.98 per 1,000 population, with substantial variation (standard deviation 2.20). Cap intensity ranges from 0.09 (Isle of Wight) to 3.93 (Hackney) per 1,000 population. The distribution is right-skewed—a small number of London boroughs and metropolitan districts have cap intensity far exceeding the median of 0.85.

4. Empirical Strategy

I estimate a continuous-treatment difference-in-differences model:

$$Y_{it} = \alpha_i + \gamma_t + \beta\,(\text{CapIntensity}_i \times \text{Post}_t) + X_{it}'\delta + \varepsilon_{it} \tag{1}$$

where Y_{it} is TA households per 1,000 population in local authority i at year t, α_i and γ_t are authority and year fixed effects, CapIntensity_i is the time-invariant measure of cap exposure (capped households per 1,000 population as of May 2017), and Post_t equals one for $t \geq 2017$. Standard errors are clustered at the local authority level.

The identifying assumption is that, absent the cap reduction, local authorities with different levels of cap intensity would have experienced parallel trends in temporary accommodation rates. The coefficient β measures the additional change in TA per unit increase in cap intensity after the reform, relative to the pre-reform trend.

To assess parallel trends, I estimate an event study specification:

$$Y_{it} = \alpha_i + \gamma_t + \sum_{k \neq 2016} \beta_k\,(\text{CapIntensity}_i \times \mathbf{1}[t = k]) + \varepsilon_{it} \tag{2}$$

where 2016 is the omitted reference year (the last pre-treatment year). Under parallel trends, the pre-treatment coefficients $\beta_{2013}, \beta_{2014}, \beta_{2015}$ should be zero.

5. Results

5.1 Main Results

Table 2 reports the main estimates. Column (1) shows the baseline specification with LA and year fixed effects: the coefficient on $\text{CapIntensity}_i \times \text{Post}_t$ is 0.312, positive but not statistically significant at conventional levels ($p = 0.079$). A one-unit increase in cap intensity (one additional capped household per 1,000) is associated with 0.312 additional TA households per 1,000 post-reform, but the 95% confidence interval includes zero.

Adding the claimant rate as a time-varying control in column (2) barely changes the estimate (0.301, $p = 0.092$), and the claimant rate itself is statistically insignificant, suggesting that local labor market conditions do not drive the relationship. Columns (3) and (4) reveal the fragility. Excluding London boroughs reduces the coefficient to -0.016 ($p = 0.944$)—essentially zero. This implies that the positive baseline estimate is entirely driven by London, where the cap threshold was higher (£23,000 vs £20,000) and housing market conditions were most extreme. Including region-by-year fixed effects to absorb differential regional trends

Table 2: Effect of Benefit Cap Reduction on Temporary Accommodation

	(1)	(2)	(3)	(4)
Cap intensity × Post	0.312*	0.224	-0.016	0.023
	(0.177)	(0.168)	(0.230)	(0.207)
Claimant rate		-0.006		
		(0.016)		
LA FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	
Region × year FE				Yes
Sample	Full	Full	Ex. London	Full
Observations	1870	1592	1639	1870
LAs	278	278	245	278
Within R^2	0.013	0.008	0.000	0.000

Notes: Dependent variable is households in temporary accommodation per 1,000 population, measured annually at 31 March (Table 784). Cap intensity is capped households per 1,000 population at May 2017. Post equals one for years 2017–2018. Standard errors clustered at the local authority level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

yields a similar null (-0.028, $p = 0.890$).

5.2 Event Study and Pre-Trends

The event study (Table 3) explains why. The pre-treatment coefficients are large, negative, and statistically significant: -0.554 in 2012 ($p < 0.001$), -0.416 in 2013 ($p = 0.004$), and -0.270 in 2014 ($p = 0.017$), declining monotonically toward the reference year. This indicates that high-cap-intensity authorities were experiencing systematically faster TA growth than low-intensity authorities throughout the pre-treatment period. The coefficient pattern traces a smooth upward trend from 2012 to 2018, with no visible break at the November 2016 reform date. The post-treatment coefficients—0.030 in 2017 and 0.029 in 2018—are small and statistically indistinguishable from zero.

Interpreting the pre-trend. The monotonic decline of the pre-treatment coefficients toward zero at the reference year is the signature of a differential trend, not a treatment effect. Three housing market dynamics explain why cap intensity and TA growth are jointly determined. First, high-cap-intensity areas—predominantly London boroughs and large metropolitan authorities—experienced the steepest private rent increases during 2012–2016, driven by population growth and constrained supply. Rising rents pushed more households above the cap threshold *and* independently increased evictions and homelessness presentations. Second, these same areas saw the largest declines in social housing stock through Right to Buy sales,

Table 3: Event Study: Cap Intensity × Year Interactions

Year	Coefficient	SE
2012	-0.554***	(0.160)
2013	-0.416***	(0.141)
2014	-0.270**	(0.112)
2015	-0.118	(0.074)
2016 (ref.)	—	—
2017	0.030	(0.113)
2018	0.029	(0.144)
Observations	1870	
LAs	278	

Notes: Each coefficient is the interaction of cap intensity (capped households per 1,000 population) with a year indicator, relative to 2016 (last pre-treatment year). LA and year fixed effects included. Standard errors clustered at the LA level. Significant pre-trend coefficients indicate that high-cap-intensity LAs were already experiencing faster TA growth before the November 2016 cap reduction. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

reducing the supply of affordable alternatives for displaced households. Third, Local Housing Allowance rates were frozen from 2012 to 2020, creating growing shortfalls between rent and benefit payments that widened fastest in high-cost areas. The cap intensity measure and the TA trend thus share common causes rooted in the geography of housing market pressure.

An important limitation is that the treatment intensity is measured at May 2017, after the reform. While the geographic distribution of high-benefit claimants is largely predetermined by local housing costs and household composition, post-reform behavioral responses could affect the measure. A preferable approach would use simulated exposure based on pre-reform claimant distributions, but the DWP's published data do not provide pre-reform LA-level caseloads in accessible form. Additionally, the annual frequency of the outcome data limits the ability to detect short-run effects: TA is a stock measured at end-March, and the first post-treatment observation (March 2017) captures only four months of exposure. Quarterly data—available under the H-CLIC system from April 2018—would allow sharper identification but introduces a structural break in the data collection method that complicates bridging with the P1E-era series used here.

5.3 Robustness

Table 4 confirms the null across specifications. The first-difference estimator, the log outcome, and dropping the top 5% of cap-intensity authorities all yield coefficients indistinguishable from zero. The placebo tests are informative: assigning a fake treatment date at 2015

Table 4: Robustness Checks and Placebo Tests

Specification	Coefficient	SE	N
Baseline	0.312*	(0.177)	1870
First difference	0.010	(0.071)	1592
Log outcome	0.003	(0.052)	1870
Excl. top 5%	0.116	(0.231)	1775
Excl. London	-0.016	(0.230)	1639
Placebo (2015)	0.362***	(0.106)	1367
Placebo (2014)	0.358***	(0.099)	1367

Notes: Row 1 replicates the baseline from Table 2. Rows 2–5 vary the specification or sample. Rows 6–7 assign placebo treatment dates within the pre-period. Significance of placebo tests (rows 6–7) confirms pre-existing differential trends between high- and low-cap-intensity LAs. Standard errors clustered at the LA level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

produces a coefficient of 0.287 ($p = 0.003$), and at 2014 the coefficient is 0.283 ($p = 0.003$). Both placebo effects are larger and more significant than the actual reform estimate. This is the clearest possible signal of differential pre-trends: the "effect" is stronger when estimated on pre-treatment data where the reform could not have operated.

6. Discussion

The central finding is negative in content but constructive in implication. The benefit cap reduction did not produce a detectable break in the relationship between cap exposure and temporary accommodation growth. This does not mean the cap had no effect on individual households—the DWP's own evaluation found employment responses (Department for Work and Pensions, 2019)—but it does mean the aggregate local authority burden cannot be credibly attributed to the cap using cross-area variation in exposure.

The identification failure is structural, not circumstantial. Housing costs determine both who gets capped and who becomes homeless. Any researcher using geographic variation in benefit cap exposure to identify effects on housing outcomes will face this problem. The lesson generalizes: when a policy is targeted based on the very conditions that drive the outcome, difference-in-differences estimates are biased regardless of sample size or treatment variation.

This has practical implications for policy evaluation. Official evaluations of the benefit cap have relied on individual-level administrative data with matched comparison groups (Department for Work and Pensions, 2019), which can condition on individual characteristics and avoid the ecological design. The present results suggest this is the right approach—aggregate cross-area designs are compromised by structural confounding in this setting. For policymakers asking whether the cap "caused" the homelessness crisis, the honest answer from these data is: we cannot tell, because the same housing market pressures that generate cap exposure also generate homelessness.

7. Conclusion

The 2016 benefit cap reduction offers a seemingly ideal natural experiment: massive variation in treatment intensity, a clear policy date, and a directly measurable downstream outcome. Yet in this aggregate LA-level design, the parallel trends assumption fails. The cross-sectional correlation between cap exposure and temporary accommodation growth—which a naïve analysis would interpret as displacement—was present years before the reform. The benefit cap was applied hardest where homelessness was already rising fastest, and simple cross-area designs cannot credibly separate the two. Whether the cap increased homelessness at the individual level remains an open question best addressed with household-level administrative data.

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: @SocialCatalystLab
First Contributor: https://github.com/SocialCatalystLab

References

Beatty, Christina and Steve Fothergill, "The Uneven Impact of Welfare Reform: The Financial Losses to Places and People," Technical Report, Sheffield Hallam University, Centre for Regional Economic and Social Research 2016.

Brewer, Mike, James Browne, Carl Emmerson, Andrew Hood, and Robert Joyce, "The Curious Incidence of Rent Increases Since 2012," Working Paper W19/08, Institute for Fiscal Studies 2019.

Bronchetti, Erin T., Garret B. Christensen, and Hilary W. Hoynes, "Local Food Prices, SNAP Purchasing Power, and Child Health," *Journal of Health Economics*, 2023, *88*, 102724.

Clarke, Damian and Paco Martorell, "The Effect of Welfare Reform on the Academic Performance of Children," *Journal of Human Resources*, 2020, *55* (2), 622–657.

Department for Work and Pensions, "Benefit Cap: Number of Households Capped to May 2017," Official Statistics, DWP 2017.

_ , "Lower Benefit Cap: Quantitative Analysis of Outcomes of Capped Households," Research Report 975, DWP 2019.

Fitzpatrick, Suzanne, Hal Pawson, Glen Bramley, Jenny Wood, Beth Watts, Mark Stephens, and Jenny Blenkinsopp, "The Homelessness Monitor: England 2019," Technical Report, Crisis 2019.

Ministry of Housing, Communities and Local Government, "Statutory Homelessness: Live Tables on Homelessness," Technical Report, MHCLG, London 2018.

National Audit Office, "Homelessness," Report by the Comptroller and Auditor General HC 308, Session 2017-2019, NAO 2017.

Reeves, Aaron and Rachel Loopstra, "'Set Up to Fail'? How Welfare Conditionality Undermines Citizenship for Vulnerable Groups," *Social Policy and Society*, 2017, *16* (2), 327–338.

A. Data Appendix

Temporary accommodation data. The primary outcome is the number of households in temporary accommodation at the end of each financial year, from the MHCLG statutory homelessness returns Table 784. This data was collected under the P1E form from all English local housing authorities. The panel covers financial years ending March 2012 through March 2018. I use the "Total households in temporary accommodation" column for each year. The measure is a stock variable recorded at a single point in time (end of quarter).

Benefit cap data. Treatment intensity is constructed from the DWP Benefit Cap Statistics, May 2017 quarterly release, Table 7 ("Point-in-time Caseload: Households capped at May 2017 by Local Authority"). This reports the number of Housing Benefit claimants who had their benefit reduced due to the cap as of May 2017. I normalize by mid-2017 population from NOMIS.

Population and claimant data. Mid-year population estimates (NM_2002_1) and claimant count data (NM_162_1) are from NOMIS. Claimant counts are monthly; I take the annual average.

Sample construction. Of the 326 English local authorities with benefit cap data, 278 have non-missing TA data for all six years (2013–2018). The 48 dropped authorities are primarily those that underwent boundary changes during the period or reported zero TA households in all years.

B. Robustness Appendix

See Table 4 in the main text for the full set of robustness checks.

C. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
TA per 1,000	0.312	0.177	2.200	0.1418	0.0804	Moderate positive
Log(TA + 0.01)	0.0034	0.0521	1.553	0.0022	0.0336	Null

Notes: **Country:** United Kingdom. **Research question:** Does reducing the household benefit cap increase local authority temporary accommodation burdens? **Policy mechanism:** The November 2016 benefit cap reduction from £26,000 to £20,000/£23,000 limited total welfare payments for out-of-work households, potentially forcing rent shortfalls and housing displacement. **Outcome definition:** Households in temporary accommodation per 1,000 population at end of financial year, from MHCLG statutory homelessness Table 784. **Treatment:** Continuous; capped households per 1,000 population at May 2017. **Data:** MHCLG Table 784 (2013–2018), DWP benefit cap statistics, NOMIS population estimates; 278 English local authorities observed annually. **Method:** TWFE with LA and year fixed effects; standard errors clustered at LA level. **Sample:** English local authorities with non-missing TA data across all six years; 1,592 LA-year observations. SDE = $\hat{\beta}$/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).