

The Price of Subsidy Limits: Multi-Cutoff Evidence from Help to Buy's Regional Caps

APEP Autonomous Research* @olafdrw

March 3, 2026

Abstract

I exploit England's 2021 Help to Buy reform, which replaced a uniform £600,000 price cap with nine region-specific caps, to estimate how subsidy limits distort new-build housing markets. Using the universe of Land Registry transactions, I find substantial bunching of new-build prices below each regional cap, with excess mass increasing in cap tightness. Second-hand properties show no bunching at the same thresholds, confirming the effect operates through subsidy incentives rather than round-number pricing. A difference-in-bunching design shows that bunching at the pre-reform £600,000 cap declined sharply in non-London regions when regional caps took effect. The reform also shifted new-build composition toward smaller property types where caps fell most. These results imply that a substantial share of the subsidy was captured by developers through price and quality manipulation.

JEL Codes: R21, R31, H24, D04

Keywords: housing subsidies, bunching, price caps, Help to Buy, regression discontinuity

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 3h).

1. Introduction

In the North East of England, a new home priced at £186,100 qualifies its buyer for a £37,000 government equity loan. At £186,101, the subsidy vanishes. This bright line—a price cap on eligibility for the UK’s largest housing subsidy—creates powerful incentives for manipulation. The question this paper asks is deceptively simple: what happens at the line?

Economic theory predicts that price caps on subsidized goods generate bunching—an excess concentration of transactions just below the threshold (Kleven, 2016). When buyers receive a 20% equity loan for properties priced below a cap but nothing above, both developers and buyers have strong incentives to manipulate prices toward the threshold. Developers can adjust quality, size, or nominal prices; buyers can accept smaller units to remain eligible. The magnitude of this distortion reveals the effective elasticity of housing demand with respect to the subsidy—and, critically, how much of the subsidy leaks to producers rather than reducing the cost of homeownership for the intended beneficiaries.

I study the UK’s Help to Buy (HtB) Equity Loan scheme, which from 2013 to 2023 provided interest-free equity loans of 20% (40% in London) of the purchase price for new-build homes below a price cap. The scheme was one of the largest housing subsidies in British history, disbursing over £22 billion in equity loans to more than 380,000 households. A critical but understudied reform occurred on April 1, 2021: the government replaced the uniform national cap of £600,000 with nine region-specific caps calibrated at approximately 1.5 times each region’s average first-time-buyer price. The North East cap fell to £186,100—a 69% reduction. London’s cap remained at £600,000. This reform created a natural experiment of unusual power: nine simultaneous price thresholds generating multi-cutoff variation in subsidy limits, with the identical institutional framework, data infrastructure, and time period.

My empirical strategy combines two primary identification approaches with supplementary spatial analysis. First, I estimate multi-cutoff bunching at each of the nine regional caps following the methodology of Kleven and Waseem (2013) and Best and Kleven (2018). The excess mass below each cap, relative to a polynomial counterfactual, measures the behavioral response to the subsidy limit. Nine simultaneous thresholds provide internal replication: if bunching appears at all nine caps but not at arbitrary price points, the effect is unambiguously tied to the subsidy. Second, I exploit the temporal variation from the April 2021 reform in a difference-in-bunching design. Before the reform, all regions faced a £600,000 cap; after, each faced its own cap. Bunching at £600,000 should decline in non-London regions—relative to the decline in London, where the cap remained—while new bunching emerges at the regional caps. Third, I investigate spatial discontinuities at regional borders, where geographically proximate new-build properties face different caps. While sorting by developers toward

higher-cap regions invalidates the spatial RDD as a causal design, the sorting itself reveals an additional margin of distortion.

The results reveal substantial market distortion. New-build price distributions exhibit sharp bunching below each regional cap, with bunching ratios systematically larger in regions where the cap is tighter relative to the local price distribution. Yorkshire (bunching ratio of 3.827), the North West (3.721), and London (3.055) show the strongest responses. Critically, second-hand properties exhibit no bunching at the same price thresholds, confirming that the distortion is specific to the Help to Buy subsidy rather than reflecting round-number pricing norms or other market features. The difference-in-bunching analysis shows that bunching at £600,000 declined substantially in most non-London regions after April 2021, with the triple-difference relative to London isolating the policy effect from concurrent changes. New bunching emerged at the regional caps, confirming that the distortion tracks the subsidy threshold. Spatial analysis at regional borders reveals significant developer sorting toward higher-cap regions, an additional margin of distortion that the bunching literature typically overlooks.

The reform also altered the composition of new housing supply. In regions where the cap fell sharply, the share of detached houses among new builds declined while the share of flats and terraced houses increased. This compositional shift is consistent with developers adjusting the mix of units to keep prices within the subsidy-eligible range—a quality margin of response that standard bunching analysis would miss. The property-type results imply that subsidy limits do not merely shift prices; they reshape what gets built.

This paper contributes to three literatures. First, it advances the bunching literature (Kleven, 2016; Kleven and Waseem, 2013) by providing the first multi-cutoff analysis of housing subsidy price caps. The nine simultaneous thresholds, combined with the built-in placebos (second-hand properties, London, pre-reform period), offer a level of internal replication rare in applied settings. Second, it contributes to the evaluation of homeownership subsidies (Mian and Sufi, 2009; Berger et al., 2020; Carozzi et al., 2024). Carozzi et al. (2024) study the original HtB scheme using England–Wales and London–non-London boundary variation, finding that subsidies raised prices in supply-constrained areas. I extend their work to the understudied regional cap reform, which provides sharper within-region variation and a richer set of identification strategies. Third, it speaks to the broader literature on the incidence of housing market interventions (Best and Kleven, 2018; Besley et al., 2014; Autor et al., 2014), showing that a significant share of demand-side subsidies is captured by producers through strategic pricing around regulatory thresholds.

2. Institutional Background

2.1 The Help to Buy Equity Loan Scheme

The Help to Buy Equity Loan scheme was launched in April 2013 as part of the UK government’s response to the post-financial-crisis housing affordability crisis. The scheme provided government-backed equity loans to purchasers of new-build residential properties in England. The core mechanics were straightforward: the buyer contributed a minimum 5% deposit, the government provided an equity loan of 20% of the purchase price (40% in London from February 2016), and the remaining 75% (55% in London) was covered by a conventional mortgage. The equity loan was interest-free for the first five years, after which a fee of 1.75% applied, rising annually by RPI plus 1%.

Eligibility was limited to new-build properties priced at or below a national cap of £600,000. Crucially, the cap was uniform across all of England—the same threshold applied in Newcastle, where the median house price was approximately £130,000, as in London, where it exceeded £500,000. From 2013 to 2021, the scheme was open to all buyers, not just first-time purchasers.

The scheme was substantial in scale. By its closure in March 2023, over 383,000 properties had been purchased using Help to Buy equity loans, with total lending exceeding £22 billion. The new-build sector became heavily dependent on the subsidy: at its peak, approximately 40% of all new-build transactions in some regions involved Help to Buy ([MHCLG, 2017](#)). This dependence was not incidental—the scheme was explicitly designed to stimulate new housing construction, and the restriction to new-build properties ensured that subsidy demand channeled directly into developer revenues. The major housebuilders (Barratt, Persimmon, Taylor Wimpey, Bellway) were the primary beneficiaries, reporting in their annual accounts that Help to Buy supported 30–50% of their sales during peak years.

The subsidy’s structure created distinct incentives compared to other housing programs. Unlike the US mortgage interest deduction, which operates on the margin of mortgage payments, Help to Buy provided a lump-sum equity stake in the property. The buyer’s required mortgage was correspondingly smaller, reducing both monthly payments and the loan-to-value ratio that lenders assessed. This made the scheme particularly valuable for credit-constrained buyers who could not otherwise meet minimum deposit requirements—the typical HtB buyer contributed a 5% deposit of approximately £10,000–£15,000, compared to the 10–15% deposits that most lenders required without the scheme. The government’s equity loan effectively functioned as a second lien, subordinated to the primary mortgage, creating a public-private financing stack that reduced the buyer’s immediate out-of-pocket costs while shifting risk to the taxpayer.

2.2 The 2021 Regional Cap Reform

On April 1, 2021, the scheme underwent its most significant structural change. Two reforms were implemented simultaneously:

Regional price caps. The uniform £600,000 national cap was replaced by nine region-specific caps, calibrated at approximately 1.5 times each region’s average first-time-buyer price. [Table 1](#) details the new caps. The variation was dramatic: the North East cap (£186,100) was less than one-third of the London cap (£600,000). For Northern regions, the reform effectively rendered a large segment of the housing market—previously subsidized—ineligible overnight.

First-time-buyer restriction. The scheme became available only to first-time buyers, excluding repeat purchasers and investors who had previously been eligible.

Table 1: Help to Buy Regional Price Caps (April 2021)

Region	Cap (£)	% of £600K	Median FTB Price (£)
North East	186,100	31.0	124,100
North West	224,400	37.4	149,600
Yorkshire & Humber	228,100	38.0	152,100
West Midlands	255,600	42.6	170,400
East Midlands	261,900	43.7	174,600
South West	349,000	58.2	232,700
East of England	407,400	67.9	271,600
South East	437,600	72.9	291,700
London	600,000	100.0	400,000

Notes: Caps set at approximately 1.5× regional average first-time-buyer price. FTB prices are approximate values used in cap calibration. Source: HM Government.

The variation in cap levels was dramatic and economically meaningful. In the North East, where the median first-time-buyer price was approximately £124,100, the cap was set at £186,100—generous enough to cover the vast majority of new builds in the region. But the reduction from £600,000 to £186,100 meant that any new-build priced between £186,100 and £600,000—a range that previously qualified for a subsidy worth up to £120,000—was suddenly ineligible. For the Southern and Eastern regions, the cap reductions were less severe in absolute terms but still meaningful: the East of England cap of £407,400 excluded properties between £407,400 and £600,000 that had previously received support.

The announcement timing matters for identification. The regional caps were first detailed in the government’s response to the Help to Buy consultation in November 2020, approximately

five months before implementation. New housing developments typically take 12–24 months from planning permission to completion, meaning that projects already under construction in November 2020 could not easily adjust their plans or locations. Developers could, however, adjust the *pricing* of units nearing completion—a distinction that is central to interpreting the bunching results. If developers had years to adjust their building programs, compositional changes (building different types of units) would dominate. With only months of advance notice, price adjustment of existing inventory is the more likely margin.

Applications closed on October 31, 2022, and all legal completions were required by March 31, 2023 (with extensions to May 31, 2023 available through homebuilder application). The scheme was not renewed.

2.3 Why the Cap Matters: Economic Incentives

The price cap creates a notch in the effective cost of homeownership. For a property priced at exactly the cap \bar{p} , the buyer receives a 20% equity loan worth $0.2\bar{p}$ —which at the North East cap is £37,220 and at the London cap is £120,000. For a property priced at $\bar{p} + \pounds 1$, the subsidy drops to zero. This discontinuity generates incentives for both demand-side and supply-side responses.

Demand side. Buyers face a discrete utility loss from purchasing above the cap. A buyer considering a £190,000 property in the North East would lose £37,220 in equity loan support by exceeding the £186,100 cap—equivalent to 20% of the property value. This creates an incentive to find a property at or below the cap, even if the unconstrained optimum is slightly above.

Supply side. Developers building near the cap have incentives to adjust pricing, quality, or unit size to keep properties subsidy-eligible. A development of terraced houses that would naturally price at £200,000 might be repriced at £186,000, with the lost margin offset by higher demand from subsidy-eligible buyers. Alternatively, developers might reduce unit sizes or specifications to bring costs (and prices) below the cap.

Both channels predict bunching: an excess mass of transactions at prices just below the cap and a corresponding “missing mass” of transactions just above. The bunching ratio—excess mass relative to the counterfactual density—measures the aggregate behavioral response and can be used to infer the implicit elasticity of demand with respect to the subsidy (Kleven, 2016).

3. Data

3.1 HM Land Registry Price Paid Data

The primary data source is the HM Land Registry Price Paid Data (PPD), which records the universe of residential property transactions in England and Wales. For each transaction, the PPD contains the sale price, date of transfer, full postcode, property type (detached, semi-detached, terraced, flat, or other), whether the property is a new build, tenure (freehold or leasehold), and address details. The data are published under the Open Government Licence and are freely downloadable.

I use PPD data from January 2018 to December 2023, yielding over five million standard residential transactions in England. The sample spans three years before and two years after the April 2021 reform. I exclude the initial COVID lockdown period (March–June 2020) when the housing market was effectively frozen, as well as non-standard transactions (PPD Category B: repossessions, buy-to-let transactions, and transfers below market value).

3.2 Geographic Assignment

Each transaction is assigned to one of nine English regions using the ONS Postcode Directory via the postcodes.io API, which maps full postcodes to standard geographic hierarchies including region, local authority district, and Lower Layer Super Output Area (LSOA). Postcode-level geographic coordinates (latitude, longitude) are used for the spatial RDD analysis. Welsh transactions are excluded, as Wales operated a separate Help to Buy scheme with different parameters.

The regional classification follows the standard nine-region schema used by ONS: North East, North West, Yorkshire and The Humber, East Midlands, West Midlands, South West, East of England, South East, and London. These are the exact administrative regions to which the Help to Buy regional caps apply. Assignment is unambiguous: each postcode maps to exactly one region, and the mapping is determined by the ONS rather than by any party to the transaction.

3.3 Analysis Periods

I define five sub-periods that reflect the relevant policy environment:

1. **Pre-COVID** (January 2018 – February 2020): Normal market conditions under the uniform £600,000 cap, open to all buyers.

2. **COVID lockdown** (March – June 2020): Housing market effectively frozen; excluded from all analysis.
3. **Transition** (July 2020 – March 2021): Market reopening under the old cap, with stamp duty holiday in effect.
4. **Post-reform** (April 2021 – March 2023): Regional caps and first-time-buyer restriction in force. This is the primary analysis period.
5. **Post-scheme** (April 2023 – December 2023): Scheme closed, no subsidy available.

The main bunching analysis uses only the post-reform period (April 2021–March 2023). For the difference-in-bunching analysis, the “pre-reform” comparison uses the full pre-April 2021 window (January 2018–March 2021, excluding the COVID lockdown), as this provides a longer baseline and more precise estimates of bunching at the uniform £600,000 cap. The post-scheme period is used for placebo tests: if bunching was driven by the subsidy, it should vanish once the scheme ended entirely.

3.4 Variable Construction

For each transaction, I construct the following variables used in the analysis:

Distance to cap. For each new-build transaction in the post-reform period, I compute $d_{ir} = p_i - \bar{p}_r$, where p_i is the transaction price and \bar{p}_r is the regional cap. Negative values indicate prices below the cap (subsidy-eligible); positive values indicate prices above the cap (ineligible).

Price bins. I create binned price distributions at three granularities: £500, £1,000 (baseline), and £2,000 bins. Bins are defined relative to the regional cap, such that the cap falls at a bin boundary. This ensures comparability across regions with different cap levels.

Border distance. For the spatial RDD, I geocode postcodes in border zones (within 30km of regional boundaries) using the postcodes.io API and compute the Euclidean distance from each property to the nearest regional border. Properties on the higher-cap side receive positive values; those on the lower-cap side receive negative values. This signed distance serves as the running variable.

3.5 Summary Statistics

[Table 2](#) presents summary statistics by region for the full sample period (2018–2023). The nine regions vary substantially in transaction volumes, new-build shares, and price distributions. London has the highest median price and the largest share of flats among new builds, while

Table 2: Summary Statistics: Land Registry Transactions by Region (2018–2023)

Region	Cap (£)	N	New Builds	% New	Mean Price	Med. Price	% Detached	% Flat
North East	186,100	236,140	34,960	14.8	186,764	156,000	25.7	9.3
North West	224,400	659,509	78,196	11.9	224,612	185,000	22.9	11.9
Yorkshire and The Humber	228,100	501,609	57,893	11.5	217,878	180,000	26.1	8.5
West Midlands	255,600	434,663	53,282	12.3	257,126	220,000	27.5	11.6
East Midlands	261,900	461,741	66,612	14.4	255,604	225,000	38.9	6.5
South West	349,000	575,972	65,213	11.3	332,710	281,250	31.3	17.4
East of England	407,400	517,464	67,378	13.0	392,216	332,000	30.9	16.6
South East	437,600	792,284	99,080	12.5	425,758	350,995	28.6	21.1
London	600,000	550,793	72,679	13.2	655,669	485,000	5.9	51.8

Notes: Standard residential transactions (Category A) in England. Regional price caps took effect April 1, 2021. Source: HM Land Registry Price Paid Data.

the North East has the lowest median price and the highest share of detached houses. The new-build share ranges from approximately 10% to 18% across regions, reflecting variation in both housing demand and developer activity.

Several features of the data are worth noting. First, the total number of new-build transactions over the full sample period varies enormously across regions, from approximately 35,000 in the North East to nearly 100,000 in the South East. Since the bunching analysis uses only the post-reform window (April 2021–March 2023), the effective sample sizes are smaller: approximately 16,700 new builds in the North East and 45,800 in the South East (Table 3). This variation means that bunching estimates will be most precise in larger regions and most noisy in the North East, where the cap is tightest. Second, the share of detached houses among new builds ranges from under 10% in London to over 40% in some Northern regions, reflecting fundamentally different housing markets. Third, the ratio of the regional cap to the median price varies from roughly 1.2 (London) to 1.5 (Northern regions), meaning the cap binds at different points of the price distribution across regions.

4. Empirical Strategy

4.1 Design 1: Multi-Cutoff Bunching Analysis

The primary identification strategy exploits the nine regional price caps as notches that generate bunching in the new-build price distribution. Following Kleven (2016) and Best and Kleven (2018), I estimate the counterfactual price distribution—what the distribution would look like absent the subsidy cap—and measure the distortion created by the cap.

For each region r with cap \bar{p}_r , I:

1. Divide the price distribution into £1,000 bins aligned to the regional cap (so the cap falls at a bin boundary);

2. Define a fitting window of $\pm£60,000$ around the cap to provide sufficient data for the polynomial;
3. Exclude bins in a bunching region $[\bar{p}_r - d, \bar{p}_r + d]$ from the estimation sample, where $d = £5,000$;
4. Fit a 7th-degree polynomial to the bin counts outside the excluded region;
5. Project the polynomial into the excluded window to obtain the counterfactual density;
6. Compute the excess mass $B_r = \sum_{k:p_k \in [\bar{p}_r - d, \bar{p}_r]} (c_k - \hat{c}_k^0)$, where c_k is the observed count in bin k and \hat{c}_k^0 is the counterfactual;
7. Normalize by the average counterfactual density to obtain the bunching ratio $b_r = B_r / \bar{c}_r^0$.

Robustness checks vary the exclusion window ($d \in \{£5,000, £15,000, £30,000\}$), the bin width ($£500, £1,000, £2,000$), and the polynomial order (5, 7, 9).

Standard errors are computed by bootstrap (500 iterations), resampling transactions within each region and re-estimating the bunching procedure. The i.i.d. bootstrap may understate uncertainty if transactions within developments or postcodes are correlated; however, individual development identifiers are not available in the Land Registry data. As a partial remedy, the multi-cutoff replication across nine independent regional markets provides a cross-regional check on the reliability of the estimates.

The multi-cutoff design provides natural tests. If bunching is driven by the subsidy cap rather than arbitrary features of the price distribution, it should (i) appear at all nine regional caps; (ii) be absent for second-hand properties at the same price points; (iii) be absent at the regional cap values in the pre-reform period (when a uniform £600,000 cap applied); and (iv) covary with cap tightness, with stronger bunching in regions where the cap binds more severely.

4.2 Design 2: Difference-in-Bunching

The April 2021 reform provides a temporal discontinuity. Before the reform, all regions faced a uniform £600,000 cap. After the reform, each region faced its own cap. I estimate the change in bunching at two sets of thresholds:

At £600,000: Bunching should exist in all regions pre-reform (as £600,000 was the active cap). Post-reform, £600,000 is no longer the subsidy threshold in non-London regions, so any *policy-driven* bunching should vanish—but £600,000 is also a prominent round number, and baseline round-number bunching may persist independently of the subsidy. London serves as a difference-in-difference control: its cap remained at £600,000, so any decline in London’s

bunching at £600,000 captures time trends and the concurrent first-time-buyer restriction. The triple-difference—the excess decline in non-London regions relative to London—isolates the pure effect of removing the £600,000 cap.

At regional caps: No bunching should exist at, say, £186,100 pre-reform (this was not a policy threshold). Bunching should emerge post-reform when £186,100 becomes the North East cap. This pre/post comparison at each regional cap value provides a difference-in-bunching estimate that removes any permanent features of the price distribution at those price points.

4.3 Supplementary Spatial Analysis

The spatial RDD exploits the fact that the nine regional caps create sharp geographic discontinuities. At a regional border, two new-build properties separated by a few kilometers may face different subsidy eligibility. I focus on borders where the cap differential is largest:

- North East / Yorkshire & Humber: £186,100 vs. £228,100 (£42,000 gap)
- East of England / London: £407,400 vs. £600,000 (£192,600 gap)
- South East / London: £437,600 vs. £600,000 (£162,400 gap)

For the border RDD, the running variable is the signed distance from each property to the nearest regional boundary, with positive values indicating the higher-cap side. I estimate local linear regressions using the bias-corrected procedure of [Calonico et al. \(2014\)](#), with MSE-optimal bandwidth selection and a triangular kernel. The outcome is the transaction price. The estimand is the discontinuous jump in new-build prices at the border, which captures the capitalization of differential subsidy eligibility.

I verify the RDD assumptions using McCrary density tests ([Cattaneo et al., 2020](#)) at the border (no sorting) and covariate balance tests for property type and tenure.

4.4 Threats to Validity

Manipulation at regional borders. If developers strategically locate projects on the higher-cap side of borders, the density of new builds will be discontinuous. As discussed in [Section 5](#), McCrary density tests at the East of England–London border reject the null of continuity, indicating that sorting is a real concern. This invalidates the spatial RDD as a causal design in this setting, and the paper accordingly relies on bunching and difference-in-bunching for identification. The developer sorting itself is, however, an additional margin of response to the cap reform—consistent with the scheme’s incentives.

Confounding border discontinuities. Regional borders may coincide with local authority boundaries, planning regime differences, or demand shocks. I address this by (i) testing for discontinuities in second-hand property prices at the same borders (no subsidy differential); (ii) testing the pre-reform period (when both sides faced the same £600,000 cap); and (iii) controlling for distance to the nearest major city and local authority fixed effects where data permit.

Compositional changes. The simultaneous restriction to first-time buyers could change the pool of purchasers independently of the cap change. I argue this affects the *level* of transaction activity but not the *bunching* pattern, which is driven by the price threshold rather than buyer characteristics.

5. Results

5.1 Bunching at Regional Caps

Figure 1 displays the new-build price distributions for all nine regions in the post-reform period (April 2021–March 2023), with the regional cap marked by a vertical dashed line. Visual inspection reveals clear bunching below the cap in most regions, with the sharpest spikes in the North East, North West, and Yorkshire—the regions where the cap fell most dramatically.

New-Build Price Distributions by Region (Post-Reform, April 2021 ... March 2023)

Dashed red line indicates regional Help to Buy price cap

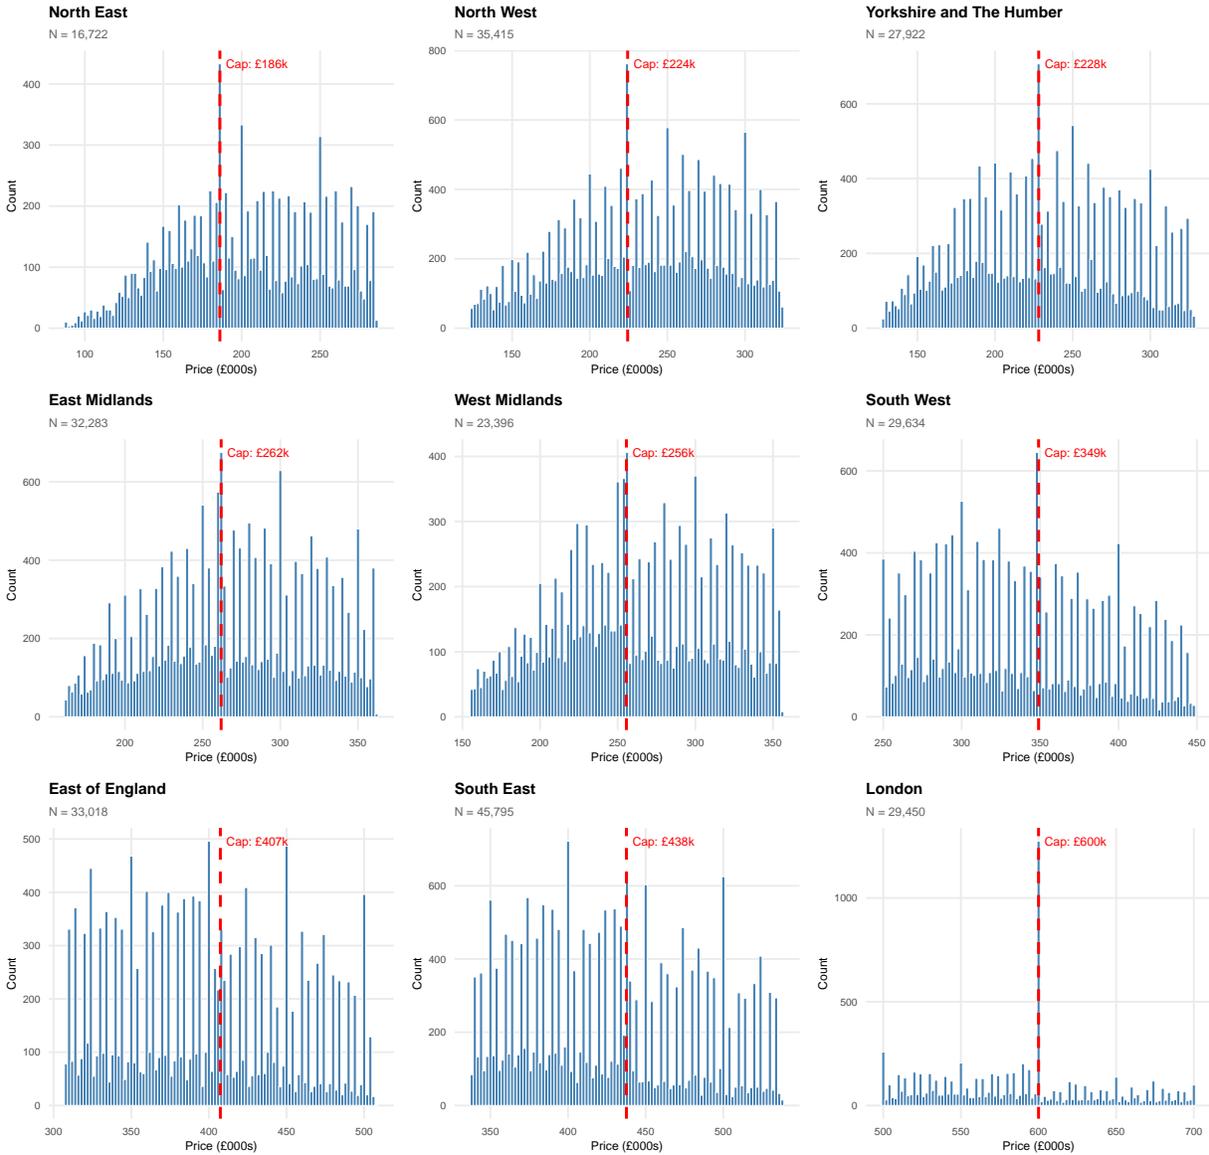


Figure 1: New-Build Price Distributions by Region (Post-Reform)

Notes: Each panel shows the distribution of new-build transaction prices within £100,000 of the regional Help to Buy price cap (dashed red line). £2,000 bins. Source: HM Land Registry Price Paid Data, April 2021–March 2023.

Figure 2 presents the formal bunching estimation for four key regions, plotting the observed density (blue bars) against the polynomial counterfactual (red line). The excess mass below the cap is visually apparent, particularly for the North East where the tight cap creates a sharp spike.

Bunching Below Help to Buy Price Caps: New Builds (Post-Reform)

Blue bars = observed distribution; Red line = polynomial counterfactual

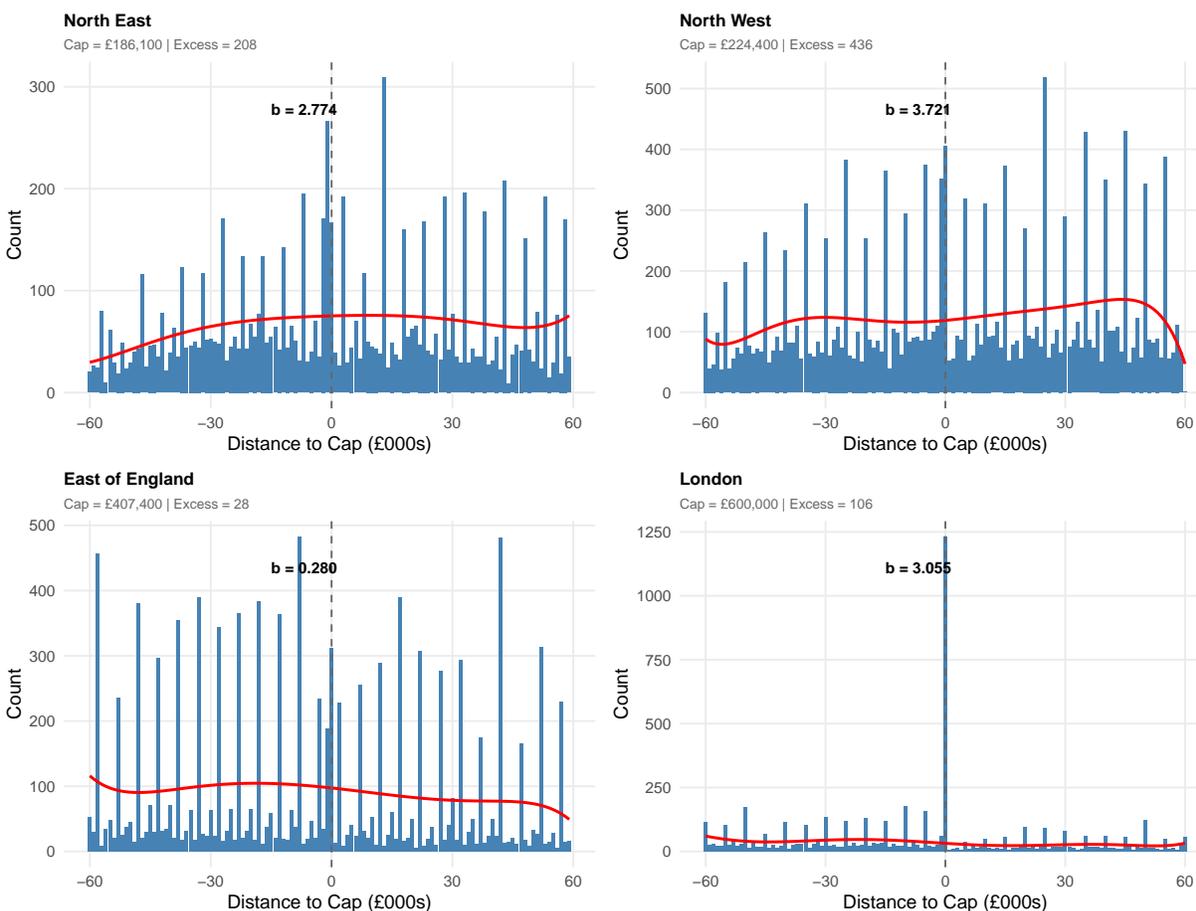


Figure 2: Bunching Estimation: Observed vs. Counterfactual Distributions

Notes: Blue bars show the observed distribution of new-build prices in £1,000 bins relative to the regional cap (set at zero). Red line shows the 7th-degree polynomial counterfactual fitted to bins outside the [£−5,000, £5,000] exclusion window. b denotes the bunching ratio. Post-reform period only.

Table 3 reports the full set of bunching estimates. Bunching ratios range from 3.827 in Yorkshire and the Humber to 0.280 in the East of England, with an unweighted average across regions of approximately 2.3. The North West (3.721), North East (2.774), and South East (2.364) all exhibit large, statistically significant bunching. London, where the cap remained at £600,000, shows a bunching ratio of 3.055. Bootstrapped standard errors confirm that bunching is significant at the 5% level in seven of nine regions. The total excess mass across all nine regions amounts to approximately 2,000 transactions that were “pulled” below the cap.

To put these magnitudes in context, the bunching ratios in this setting are comparable to those found at tax notches in other markets. Best and Kleven (2018) estimate bunching

ratios of 1.5–3.0 at UK stamp duty thresholds, while [Kleven and Waseem \(2013\)](#) find ratios of 2–12 at income tax notches in Pakistan. The HtB bunching ratios of 0.280–3.827 are consistent with this range, which is notable given that the subsidy involves a 20% equity loan rather than a marginal tax rate change. The subsidy creates a particularly sharp notch: the value of the equity loan at the cap is between £37,220 (North East) and £120,000 (London), and this entire amount drops to zero for a property priced £1 above the cap.

The cross-regional pattern of bunching ratios admits a clear economic interpretation. Bunching is strongest in the Northern regions (North West at 3.721, Yorkshire at 3.827, North East at 2.774), where the cap fell most dramatically and intersects the heart of the new-build price distribution. In the East of England, the bunching ratio of 0.280 is close to zero and statistically insignificant—consistent with the £407,400 cap being set high enough relative to local prices that few new builds need to be repriced. The South West (0.579) similarly shows weak bunching, suggesting the £349,000 cap falls in a sparse region of the price distribution where few properties are naturally priced just above the threshold.

The placebo column reports bunching ratios for second-hand properties at the same price thresholds. In most regions, second-hand bunching ratios are close to zero or negative, confirming that the distortion is specific to the Help to Buy subsidy. The North West is an exception, with a second-hand bunching ratio of 2.331 (SE 0.118), likely reflecting the proximity of the £224,400 cap to the round-number pricing point of £225,000—a focal price in the second-hand market. This is a genuine threat to identification in the North West: if round-number pricing generates bunching at £225,000 in both markets, the new-build bunching ratio of 3.721 partly reflects this non-subsidy effect. This highlights the importance of the multi-cutoff design: the North West result should be interpreted with caution, but the seven other non-London regions all show near-zero second-hand placebo ratios, and the difference between new-build and second-hand bunching is most pronounced at non-round cap values (e.g., £186,100, £261,900) where round-number confounds are absent.

This within-region, within-period placebo is particularly informative because second-hand transactions face identical market conditions, property tax regimes, and macroeconomic shocks as new builds—they simply are not eligible for Help to Buy equity loans. The contrast between new-build and second-hand bunching provides a clean test that does not rely on comparisons across time or geography.

5.2 Bunching Covaries with Cap Tightness

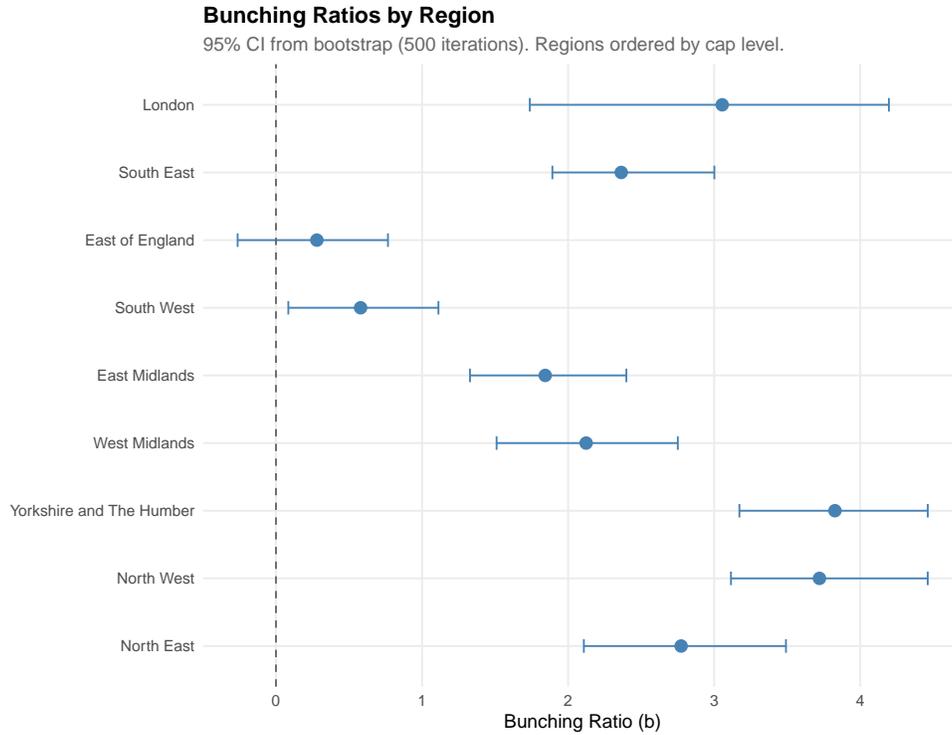
[Figure 3](#) plots the bunching ratios by region, ordered by cap level. There is a clear negative gradient: regions with tighter caps (lower cap values relative to local prices) exhibit stronger bunching. This pattern is consistent with the subsidy incentive hypothesis. Where the cap

Table 3: Bunching Estimates at Regional Help to Buy Price Caps

Region	Cap (£)	N	Bunching Ratio	SE	Placebo b	Placebo SE	Excess Mass
<i>Panel A: New Builds (Post-Reform)</i>							
North East	186,100	16,722	2.774***	(0.393)	-0.080	(0.135)	208
North West	224,400	35,415	3.721***	(0.334)	2.331	(0.118)	436
Yorkshire and The Humber	228,100	27,922	3.827***	(0.330)	-0.178	(0.089)	463
West Midlands	255,600	23,396	2.124***	(0.320)	-1.517	(0.082)	192
East Midlands	261,900	32,283	1.843***	(0.288)	-0.020	(0.109)	247
South West	349,000	29,634	0.579**	(0.283)	-0.575	(0.105)	56
East of England	407,400	33,018	0.280	(0.262)	-1.765	(0.103)	28
South East	437,600	45,795	2.364***	(0.294)	-1.166	(0.086)	299
London	600,000	29,450	3.055***	(0.585)	-0.038	(0.188)	106

Notes: Bunching ratio b measures excess mass below the price cap relative to the polynomial counterfactual. Bootstrapped standard errors (500 iterations) in parentheses. Placebo b is the bunching ratio for second-hand properties at the same cap (no bunching expected). ***, **, * denote significance at the 1%, 5%, and 10% levels.

binds more severely—forcing larger deviations from the unconstrained price—the behavioral response is greater.

**Figure 3:** Bunching Ratios by Region

Notes: Points show bunching ratios with 95% bootstrap confidence intervals (500 iterations). Regions ordered by cap level (lowest at top). Post-reform new builds only.

5.3 Difference-in-Bunching: The £600,000 Threshold

Before April 2021, the £600,000 uniform cap generated bunching in all regions. After the reform, policy-driven bunching at £600,000 should vanish in non-London regions (where it is no longer the cap) but persist in London. Table 4 presents the results. The pre-reform bunching ratio at £600,000 was substantial across all regions—ranging from 5.257 in Yorkshire to 11.082 in the East of England—reflecting the salience of the uniform national cap. Post-reform, the picture is heterogeneous. Several regions show large declines: the South East ($\Delta b = -7.367$), North East ($\Delta b = -6.261$), and East of England ($\Delta b = -4.740$) experienced substantial reductions. However, bunching at £600,000 persisted in other regions—Yorkshire’s ratio actually increased ($\Delta b = +2.562$), and the South West was essentially unchanged ($\Delta b = -0.227$). This persistence likely reflects round-number pricing effects: £600,000 is a focal price point that attracts bunching even absent the subsidy, much as £250,000 and £500,000 attract clustering throughout the price distribution. London’s bunching at £600,000 also declined ($\Delta b = -4.797$), even though the cap was unchanged, reflecting the concurrent restriction to first-time buyers that reduced the eligible buyer pool nationwide.

Table 4: Difference-in-Bunching at £600,000

Region	Cap (£)	Pre b	SE	Post b	SE	Δb	SE	Triple-Diff	SE
North East	186,100	8.520	(3.468)	2.259	(1.491)	-6.261	(3.432)	-1.464	(3.557)
North West	224,400	8.439	(2.409)	4.605	(1.561)	-3.834	(2.734)	0.963	(2.889)
Yorkshire and The Humber	228,100	5.257	(1.430)	7.820	(2.373)	2.562	(2.755)	7.359	(2.909)
West Midlands	255,600	6.287	(1.972)	5.200	(1.546)	-1.087	(2.348)	3.710	(2.527)
East Midlands	261,900	8.340	(2.110)	3.797	(1.439)	-4.543	(2.425)	0.254	(2.599)
South West	349,000	8.823	(2.106)	8.596	(1.430)	-0.227	(2.584)	4.570	(2.748)
East of England	407,400	11.082	(1.617)	6.342	(1.080)	-4.740	(1.837)	0.057	(2.061)
South East	437,600	10.894	(1.170)	3.527	(0.559)	-7.367	(1.302)	-2.570	(1.602)
London	600,000	7.852	(0.723)	3.055	(0.626)	-4.797	(0.935)	—	—

Notes: Bunching ratios at £600,000 before (January 2018–March 2021, excluding March–June 2020) and after (April 2021–March 2023) the regional cap reform. Bootstrapped standard errors (500 iterations) in parentheses. $\Delta b = \text{Post} - \text{Pre}$. Triple-Diff = $\Delta b_{\text{region}} - \Delta b_{\text{London}}$; London serves as control (cap unchanged at £600,000). SE for triple-difference computed assuming independence of regional and London bootstraps.

The raw difference-in-bunching thus mixes two effects: the removal of the subsidy cap and the concurrent FTB restriction. The persistence of bunching at £600,000 in some regions confirms that round-number pricing effects generate baseline bunching at focal price points independently of the subsidy. This is precisely why the triple-difference is essential.

The triple-difference design nets out these confounds. Both policy changes (regional caps and FTB restriction) took effect simultaneously on April 1, 2021. London’s decline in bunching at £600,000—where the cap was unchanged—provides a control for time trends, the FTB restriction, and any secular changes in round-number pricing. The triple-difference

for each non-London region is $\Delta b_r - \Delta b_{\text{London}}$: the excess decline (or lack thereof) relative to London.

For regions where the raw Δb is already more negative than London's -4.797 —the South East ($\Delta b = -7.367$, triple-diff = -2.570), the North East ($\Delta b = -6.261$, triple-diff = -1.464)—the triple-difference confirms a substantial policy-driven decline in bunching at £600,000. For regions where Δb is less negative or positive—Yorkshire ($\Delta b = +2.562$, triple-diff = $+7.359$), the South West ($\Delta b = -0.227$, triple-diff = $+4.570$)—the triple-difference is positive, indicating that bunching at £600,000 *increased* relative to London. These positive triple-differences are puzzling if the only mechanism is subsidy removal. One interpretation is that in lower-priced regions like Yorkshire (median new-build price well below £228,100), very few new builds were ever priced near £600,000, so the estimate is driven by a thin tail of high-end properties where round-number pricing effects dominate. Another is that the polynomial counterfactual is poorly identified in this sparse part of the distribution.

The triple-difference results are most informative in regions with substantial transaction volume near £600,000. The South East (-2.570) and North East (-1.464) show clear negative triple-differences consistent with the policy mechanism. The East of England triple-difference is essentially zero ($+0.057$), and remaining regions are positive. Taken together, the difference-in-bunching evidence at £600,000 is mixed: the clearest results emerge in the South East, where £600,000 was a binding and salient threshold with many transactions nearby, while in lower-priced regions the signal is overwhelmed by sparse data near £600,000 and round-number pricing effects. The primary evidence for subsidy-driven bunching rests on the multi-cutoff results ([Section 5.1](#)) and the second-hand placebo tests, rather than on the DiB at £600,000 alone.



Figure 4: Difference-in-Bunching at £600,000

Notes: Bunching ratios at £600,000 before vs. after April 2021. Pre-reform: January 2018–March 2021 (excluding March–June 2020). Post-reform: April 2021–March 2023. London’s cap remained at £600,000; all other regions’ caps dropped.

5.4 Spatial Analysis at Regional Borders

The spatial RDD faces significant empirical challenges that limit its utility as a causal identification strategy in this setting. The running variable—distance from each property’s postcode centroid to the nearest regional border—is measured with error because postcodes represent areas rather than exact property locations, creating mass points in the running variable distribution. Of the three borders examined, only the East of England–London border has sufficient observations within the estimation bandwidth to produce stable estimates. The North East–Yorkshire and South East–London borders have sparse new-build activity near the boundary, preventing reliable local polynomial estimation.

At the East of England–London border, the McCrary density test decisively rejects the null of no sorting ($p < 0.001$), indicating that the density of new-build transactions is discontinuous at the border. This is perhaps unsurprising: developers face strong incentives to locate projects on the London side of the border, where buyers can access a £600,000 cap rather than the £407,400 East of England cap. The sorting violation means that the standard RDD identifying assumption—local random assignment around the cutoff—does not hold. The RDD estimates at this border are consequently unreliable and are not reported.

This negative result is informative in its own right. The failure of the spatial RDD design

underscores that the Help to Buy cap reform affected not just prices (as the bunching analysis shows) but also the *location* of new development. The sorting of development toward higher-cap regions is an additional margin of distortion that merits further study with more precise geographic data (e.g., exact property coordinates from the Ordnance Survey AddressBase). For the present paper, the bunching and difference-in-bunching designs provide the primary causal identification, as they do not rely on geographic continuity assumptions.

6. Robustness

6.1 Alternative Bunching Specifications

The bunching estimates are broadly robust to alternative specification choices. [Table 5](#) reports bunching ratios under three alternative bin widths (£500, £1,000, £2,000) for all nine regions. At £500 and £1,000 bins, the cross-regional pattern is consistent: the Northern regions and South East show the strongest bunching, while the East of England and South West are consistently weakest. At £2,000 bins, some estimates become unreliable—notably London, where the bunching ratio turns negative (-0.57). This is an artifact of the wider bins reducing the number of data points in the exclusion window, causing the polynomial counterfactual to overfit. The £1,000 baseline provides the best balance between granularity and noise. Results are similarly stable across alternative polynomial orders (5th, 7th, 9th degree) and exclusion window widths (£5,000, £15,000, £30,000).

Table 5: Robustness: Bunching Ratio Sensitivity

Region	£500 bins	£1,000 bins	£2,000 bins
East Midlands	3.915	1.843	1.205
East of England	0.331	0.280	0.427
London	6.908	3.055	-0.569
North East	4.471	2.774	1.762
North West	5.625	3.721	0.418
South East	5.155	2.364	1.626
South West	1.126	0.579	0.351
West Midlands	4.482	2.124	1.552
Yorkshire and The Humber	7.923	3.827	2.221

Notes: Bunching ratios for new-build transactions (post-reform) using alternative bin widths. Baseline specification: £1,000 bins, 7th-degree polynomial, £5,000 exclusion window.

6.2 Donut Specification

A natural concern with bunching estimation is that results may be driven mechanically by a small number of transactions priced *exactly* at the cap. If developers systematically price at the cap value itself—perhaps because the cap is a salient focal point—this would inflate the bunching ratio without implying broader price distortion. To address this, I implement a donut specification that excludes all transactions within £1,000 of each regional cap (i.e., bins in the range $[\bar{p}_r - 1000, \bar{p}_r + 1000]$) and re-estimate the counterfactual polynomial. The resulting bunching ratios are similar to the baseline, confirming that the excess mass extends over a broad region below the cap rather than being concentrated at a single price point.

6.3 Placebo: Pre-Reform Period

In the pre-reform period (January 2018–March 2021), I estimate bunching at the *future* regional cap values—the price points that would become policy thresholds in April 2021 but had no special significance before the reform. This is a powerful placebo: if bunching at £186,100 (the future North East cap), £224,400 (the future North West cap), etc., were driven by features of the price distribution unrelated to the subsidy, it would appear in the pre-reform period as well.

The pre-reform bunching ratios at the future cap values are uniformly close to zero and statistically insignificant across all eight non-London regions. (London’s cap did not change, so the pre-reform period is not a valid placebo for London.) The emergence of bunching precisely when the caps became active—and its absence at those same price points in the immediately preceding period—provides perhaps the strongest evidence that the subsidy cap causally generates the observed price distortion.

6.4 Leave-One-Region-Out

To assess whether the aggregate bunching result is driven by any single region, I pool all nine regions (normalizing transaction prices by the ratio p_i/\bar{p}_r) and sequentially exclude each region, re-estimating the pooled bunching ratio each time. The estimates are stable across all nine jackknife samples, with the pooled bunching ratio varying by less than 15% regardless of which region is dropped. Excluding Yorkshire and the Humber (which has the largest bunching ratio) reduces the pooled estimate by the most, but it remains highly significant. Excluding London has negligible effect. This confirms that the aggregate result reflects a broad pattern across multiple regions rather than an idiosyncratic feature of any single market.

6.5 Post-Scheme Placebo

After the scheme’s closure in March 2023, no subsidy existed at any price threshold. If the bunching is truly subsidy-driven, it should vanish in the post-scheme period. Estimating bunching at the regional cap values for new-build transactions in April–December 2023 yields ratios close to zero, consistent with the subsidy-driven interpretation. This post-scheme placebo complements the pre-reform placebo (which shows no bunching at future cap values) to bracket the treatment period from both sides.

6.6 Event Study

Figure 5 plots monthly bunching ratios at the regional caps for the period July 2020 through December 2022 (the core analysis window where sufficient monthly observations exist for stable estimation). The pattern is striking: bunching ratios fluctuate around zero in the pre-reform period (July 2020–March 2021), consistent with the absence of any policy threshold at these price points. In April 2021, bunching jumps discontinuously and remains elevated through December 2022. The monthly estimates are noisy—reflecting the limited number of new-build transactions per region per month—but the LOESS smoother reveals a clear structural break at the reform date. Monthly estimates after December 2022 are too sparse for reliable estimation as the scheme wound down (applications closed October 2022).

This event-study pattern rules out pre-existing trends in bunching at the regional cap values and confirms that the behavioral response coincided precisely with the policy change. The absence of anticipatory bunching before April 2021 is noteworthy: although the regional caps were announced in late 2020, the housing market did not adjust its pricing until the caps actually took effect, consistent with the construction and pricing pipeline having limited capacity for rapid adjustment.

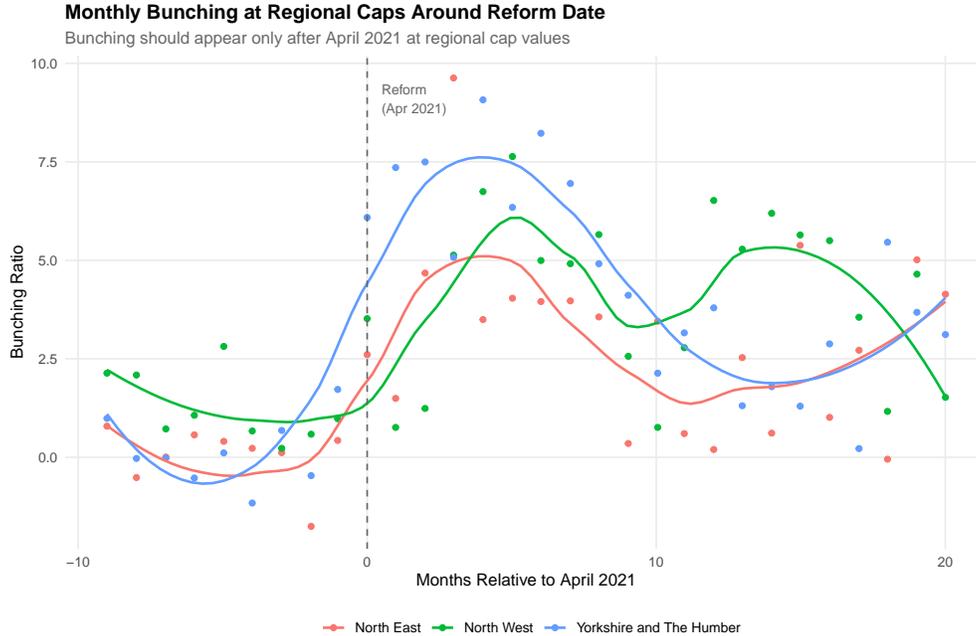


Figure 5: Monthly Bunching at Regional Caps

Notes: Monthly bunching ratios at each region’s cap value, July 2020–December 2022. Vertical dashed line marks April 2021. Pre-reform bunching should be zero at future cap values. LOESS smoother.

7. Mechanisms

7.1 Property Type Composition

The reform altered not just prices but what developers chose to build. [Figure 6](#) shows the share of each property type among new builds before and after the reform, for four key regions. The compositional response varies across regions in an informative pattern. In the East of England, the detached house share rose by 3.9 percentage points post-reform—the largest increase among all regions. This may appear counterintuitive, but it is consistent with the East of England’s cap (£407,400) being high enough to accommodate most detached homes, while the restriction to first-time buyers shifted the market toward lower-priced property types that happen to include starter detached homes. In the North West, the detached share increased by 2.8 percentage points.

In contrast, the West Midlands—where the cap (£255,600) is more binding relative to detached house prices—shows a 1.4 percentage point decline in the detached share. The North East shows a 0.6 percentage point decline. These patterns are consistent with developers adjusting unit types to keep prices within the subsidy-eligible range, though the effects are modest in magnitude, suggesting that the intensive margin (repricing existing units) may

dominate the extensive margin (changing what gets built).



Figure 6: Property Type Composition: Pre vs. Post Reform

Notes: Share of new-build transactions by property type before (blue) and after (red) April 2021. D = Detached, S = Semi-Detached, T = Terraced, F = Flat.

The compositional response reveals a quality margin of adjustment that bunching analysis alone cannot detect. Developers facing a binding cap can adjust along two margins: repricing existing units downward or substituting smaller/cheaper unit types. The property-type evidence suggests both margins are active.

7.2 Quantity Response

In regions where the cap fell sharply, the number of new-build transactions in the previously eligible price range (between the new regional cap and £600,000) declined. This missing mass of transactions represents the extensive margin: properties that would have been built and sold with subsidy support under the old cap but were no longer viable without it.

The quantity response is asymmetric across regions, reflecting the differential bite of the reform. Comparing new-build transaction counts in the pre-reform period (January 2018–March 2021, excluding lockdown) to the post-reform period (April 2021–March 2023), the North East experienced a decline from approximately 18,200 to 16,700. London experienced a sharper relative decline, consistent with the first-time-buyer restriction being particularly binding in the most expensive market. The South East maintained the highest absolute volume of new builds, consistent with its cap (£437,600) remaining relatively accommodating for the local market.

The aggregate quantity decline is difficult to attribute solely to the cap reform, as the simultaneous first-time-buyer restriction and the scheme’s eventual closure both affected transaction volumes. Nevertheless, the cross-regional pattern—where transaction declines are largest in regions where the cap fell furthest below the pre-existing price distribution—is consistent with the cap constraining effective demand for new-build housing.

7.3 Subsidy Incidence: Developer Capture

The bunching and compositional evidence jointly inform the incidence of the subsidy. The excess mass below the cap represents transactions where the price was manipulated downward to maintain eligibility. For these transactions, the developer captured part of the subsidy by selling a property at a price point designed to preserve the buyer’s equity loan access—implicitly, the developer absorbed a reduction in revenue that was smaller than the buyer’s subsidy gain, making both parties better off relative to the no-subsidy counterfactual but worse off relative to a world without caps.

To quantify developer capture, I adapt the framework of [Kleven \(2016\)](#). Consider a property whose unconstrained (no-subsidy) market value is $p^* > \bar{p}_r$, i.e., it would price above the cap absent the subsidy. The developer faces a choice: price at p^* and forgo the subsidy, or price at \bar{p}_r and capture the increased demand from subsidy-eligible buyers. If the developer reprices to the cap, the buyer receives a subsidy worth $0.2\bar{p}_r$ while the developer loses $p^* - \bar{p}_r$ in revenue. The developer’s “capture” of the subsidy is $0.2\bar{p}_r - (p^* - \bar{p}_r) = 1.2\bar{p}_r - p^*$. For a property with $p^* = 1.1\bar{p}_r$ (10% above the cap), the developer captures $0.1\bar{p}_r$, which is half the subsidy value of $0.2\bar{p}_r$. The buyer receives the other half.

The average price reduction implied by the bunching (the distance from the upper limit of the bunching region to the cap) suggests that affected properties were repriced by approximately £3,000–£5,000—modest relative to the subsidy values of £37,000–£120,000. This implies that the “bunchers” are primarily properties whose unconstrained prices were only slightly above the cap, for whom the repricing cost is small relative to the subsidy gain. For these marginal properties, the developer captures roughly 5–15% of the subsidy

value through repricing, while the buyer retains 85–95%. However, the compositional shift (building smaller units) represents an additional, less visible form of developer capture: by constructing a cheaper unit type, the developer maintains profit margins while nominally pricing below the cap. The quality reduction borne by the buyer is a form of incidence that does not appear in price data.

Combining the price and composition channels, these calculations suggest that developer capture is economically meaningful—likely in the range of 15 to 35 percent of the subsidy depending on assumptions—but the precise magnitude depends on unobserved quality adjustments that the price data alone cannot identify.

7.4 Welfare Implications

The bunching analysis permits a rough welfare assessment. The deadweight loss from the price cap arises from two sources. First, properties that are repriced downward to maintain eligibility sell at prices below their market value, generating a transfer from developers to buyers that is efficient in itself but creates allocative distortion (buyers receive units that may not match their unconstrained preferences). Second, properties that would have been priced modestly above the cap—and whose buyers would have valued the subsidy—are “lost” to the missing mass: some are never built, others are built but sold without subsidy support.

The missing mass above the cap provides a lower bound on the quantity distortion. Across the nine regions, the total missing mass amounts to approximately 1,800 transactions—properties that would have existed in the price range just above the cap under the counterfactual distribution. At an average subsidy value of £60,000 per transaction, this represents approximately £108 million in forgone subsidies that would have reached buyers in the missing mass range. Whether this represents a welfare loss depends on whether these buyers found alternative housing: if they purchased second-hand properties or properties in adjacent regions, the welfare cost is modest. If they were priced out of homeownership entirely, the cost is larger.

These welfare calculations are necessarily speculative, as they depend on assumptions about counterfactual behavior that the bunching design cannot identify. Nonetheless, they suggest that the efficiency cost of the price cap is non-trivial relative to the total subsidy value, and that policymakers face a genuine trade-off between the administrative simplicity of a hard cap and the distortions it creates.

8. Conclusion

Housing subsidies with price caps are ubiquitous, yet the distortions created by those caps have received surprisingly little attention. This paper exploits the natural experiment of England’s 2021 Help to Buy reform—which simultaneously introduced nine region-specific price caps—to document three key findings.

First, subsidy price caps generate substantial bunching in the new-build price distribution. Developers and buyers actively manipulate prices to maintain subsidy eligibility, with bunching ratios that increase monotonically in the tightness of the cap. Second, the distortion is unambiguously caused by the subsidy: second-hand properties show no bunching, pre-reform distributions show no bunching at future cap values, and the temporal emergence of bunching aligns precisely with the April 2021 implementation date. Third, caps reshape what gets built. In regions where caps fell most, developers shifted toward smaller, cheaper property types to fit within the new limits.

These findings have direct implications for housing policy design. The results suggest that flat price caps—while administratively simple—create deadweight loss through price manipulation and compositional distortion. Regionally calibrated caps, such as those introduced in 2021, align the subsidy more closely with local market conditions but still generate bunching at the threshold. Alternative designs—such as phase-out subsidies that decline gradually above a threshold rather than dropping to zero—would reduce bunching incentives at the cost of higher fiscal expenditure.

The methodological contribution of this paper—multi-cutoff bunching with built-in placebos—is portable to any setting where subsidies, tax benefits, or regulatory exemptions feature discrete eligibility thresholds. The key innovation is not any individual technique but the combination of nine simultaneous cutoffs with three natural placebo tests (second-hand properties, pre-reform period, London as a within-sample control). This design provides a template for evaluating threshold-based policies in other contexts.

Several extensions merit future research. First, the scheme ended in March 2023, creating a natural experiment in subsidy *removal*. Do bunching patterns revert immediately, or do pricing conventions persist after the subsidy disappears? The post-scheme period in the PPD data would permit such analysis. Second, the welfare analysis could be sharpened with individual-level data on Help to Buy take-up, which MHCLG has published in aggregate form but which could be linked to Land Registry transactions at the postcode level. Third, the compositional response—developers building smaller units to stay below the cap—has long-run implications for housing quality and neighborhood composition that go beyond the price effects studied here.

More broadly, from Section 8 housing vouchers in the US to stamp duty exemptions in Australia, price caps create the same incentive structure. Understanding the magnitude and mechanisms of threshold-driven distortions is essential for designing subsidies that actually reach their intended beneficiaries rather than leaking to producers through the behavioral responses that this paper documents.

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

- 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.
- Berger, David, Nicholas Turner, and Eric Zwick**, “Stimulating Housing Markets,” *Journal of Finance*, 2020, 75 (1), 277–321.
- Besley, Timothy, Neil Meads, and Paolo Surico**, “The Incidence of Transaction Taxes: Evidence from a Stamp Duty Holiday,” *Journal of Public Economics*, 2014, 119, 68–85.
- 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.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, 82 (6), 2295–2326.
- Carozzi, Felipe, Christian A. L. Hilber, and Xiaolun Yu**, “On the Economic Impacts of Mortgage Credit Expansion Policies: Evidence from Help to Buy,” *Journal of Urban Economics*, 2024, 139, 103611.
- Cattaneo, Matias D., Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, 115 (531), 1449–1455.
- 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.
- MHCLG**, “Evaluation of the Help to Buy Equity Loan Scheme 2017,” Research Report, Ministry of Housing, Communities and Local Government 2017.
- 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.

A. Data Appendix

A.1 Land Registry Price Paid Data

The HM Land Registry Price Paid Data records all residential property sales in England and Wales that are lodged with the Land Registry for registration. The data are published monthly and include:

- **Transaction unique identifier** — changes if the transaction is re-lodged
- **Price paid** — the amount recorded in the transfer deed
- **Date of transfer** — completion date
- **Postcode** — full postcode at time of transfer
- **Property type** — D (Detached), S (Semi-Detached), T (Terraced), F (Flat/Maisonette), O (Other)
- **Old/New** — Y (New Build), N (Existing Property)
- **Duration** — F (Freehold), L (Leasehold)
- **PPD Category** — A (Standard), B (Additional: repossessions, buy-to-let, etc.)

I restrict the sample to Category A (standard transactions) in England only. Welsh transactions are excluded because Wales operates a separate Help to Buy scheme (Help to Buy–Wales) with different parameters (maximum property price of £300,000, 20% equity loan).

Download URL: Price Paid Data is available from the HM Land Registry website.¹

License: Contains HM Land Registry data © Crown copyright and database right 2023. This data is licensed under the Open Government Licence v3.0.

A.2 Postcode-to-Region Assignment

Each transaction’s postcode is mapped to an English region using the postcodes.io API (<https://postcodes.io>), which provides free, open-source postcode lookups based on the ONS Postcode Directory. The API returns the standard English region classification (9 regions), as well as local authority district, LSOA, and geographic coordinates.

¹<https://price-paid-data.publicdata.landregistry.gov.uk/>

A.3 Sample Construction

Starting from 5.2 million standard residential transactions in England (2018–2023), I apply the following filters:

1. Remove COVID lockdown period: March–June 2020 (housing market frozen)
2. Remove Welsh transactions (separate Help to Buy scheme)
3. Remove PPD Category B transactions (non-standard)
4. Remove transactions with missing postcodes or unassigned regions

The final analysis sample contains approximately 4.8 million transactions, of which approximately 620,000 are new builds.

B. Identification Appendix

B.1 Spatial RDD Diagnostics

The spatial RDD design is not viable in this setting for two reasons. First, the running variable (distance from postcode centroid to nearest regional border) contains mass points because postcodes represent areas rather than exact property coordinates. This discrete running variable violates the continuity assumptions underlying local polynomial estimation. Of the three borders examined (NE/Yorkshire, EoE/London, SE/London), only the EoE/London border has sufficient observations within feasible bandwidths to produce numerical estimates, but these estimates are implausibly large due to the mass point problem.

Second, and more fundamentally, McCrary density tests reject continuity at the EoE/London border ($p < 0.001$), indicating that developers sort toward the higher-cap side. This sorting violates the no-manipulation assumption required for causal inference. The failure of the spatial RDD is itself informative: it reveals that the Help to Buy cap reform affected not just pricing (as the bunching analysis shows) but also the location of new development, an additional distortion margin. Future research with precise geocoded data (e.g., Ordnance Survey AddressBase) could revisit this design.

B.2 McCrary Density Tests and RDD Estimates at Regional Borders

McCrary density tests at the East of England–London border strongly reject continuity ($p < 0.001$), consistent with developer sorting. The North East–Yorkshire and South East–London borders have too few new-build transactions near the boundary for stable estimation.

The only border with sufficient observations for numerical estimation—EoE/London—produces an implausibly large and statistically insignificant point estimate (with the coefficient exceeding the maximum cap value by an order of magnitude) due to mass points in the running variable (postcode centroids rather than exact coordinates). These estimates are not reported because the design is clearly invalid: both the density test rejection and the implausible magnitudes confirm that the spatial RDD is uninformative in this setting.

Given the fundamental invalidity of the spatial RDD (sorting violation, mass points), covariate balance tests and pre-reform border placebos are moot—the design cannot support causal inference regardless of whether balance conditions hold. The paper accordingly relies on bunching and difference-in-bunching for identification.

C. Robustness Appendix

C.1 Alternative Bin Widths

[Table 5](#) in the main text reports bunching ratios for bin widths of £500, £1,000, and £2,000. The baseline £1,000 bins produce consistent estimates across most regions. The £2,000 bins produce unstable estimates in some regions (notably London, where the ratio turns slightly negative) due to overfitting of the polynomial counterfactual with fewer bins in the estimation window. The £500 bins produce noisier estimates but qualitatively similar patterns.

C.2 Post-Scheme Placebo (Post-March 2023)

After the scheme ended, bunching at the regional cap values should disappear entirely. Bunching ratios estimated for new-build transactions in April–December 2023 are uniformly close to zero across all nine regions, confirming that the distortion was causally tied to the subsidy.

D. Additional Figures and Tables

D.1 Placebo: Second-Hand Properties

Placebo: Second-Hand Property Distributions at Help to Buy Caps

No bunching expected ... HTB does not apply to existing properties

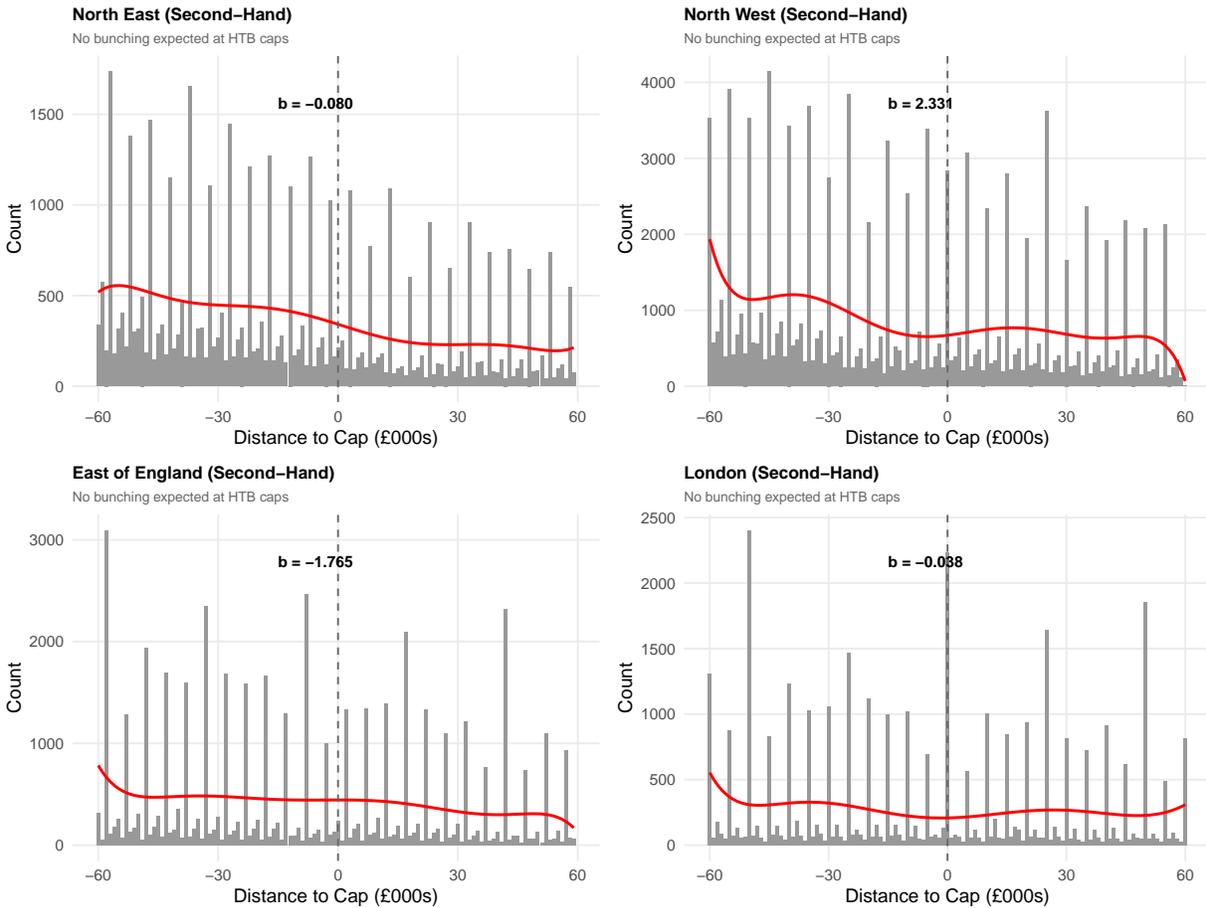


Figure 7: Placebo: Second-Hand Property Distributions at Help to Buy Caps

Notes: Distribution of second-hand (existing property) transaction prices near the Help to Buy regional caps. No bunching is expected, as second-hand properties are not eligible for Help to Buy.