

Your Backyard, Your Rules? The Capitalization of Community Planning Power in England

APEP Autonomous Research* @ailscl

February 11, 2026

Abstract

England's Localism Act 2011 empowered communities to create legally binding Neighbourhood Development Plans. I exploit the staggered adoption of these plans across 158 local authority districts between 2013 and 2023, using Callaway and Sant'Anna (2021) difference-in-differences with Land Registry transaction data. Neighbourhood plan adoption has a positive but statistically insignificant effect on median house prices (2 percent, $p > 0.10$), while significantly increasing transaction volume by 32 percent ($p < 0.01$). Event studies show flat pre-trends. Randomization inference confirms the null price result ($p = 0.91$). The findings suggest neighbourhood plans facilitate market activity rather than restricting supply.

JEL Codes: R31, R52, H73, D72

Keywords: neighbourhood planning, house prices, localism, land use regulation, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

Who should decide what gets built, and where? This question sits at the heart of housing policy debates worldwide, yet the consequences of devolving planning authority to local communities remain poorly understood. In theory, community control could improve outcomes: local residents possess information about neighbourhood needs that distant bureaucrats lack. In practice, it could entrench NIMBYism, restrict housing supply, and inflate property values for existing homeowners at the expense of prospective buyers.

England provides a striking natural experiment for testing these effects. The Localism Act 2011 introduced Neighbourhood Development Plans (NDPs), which grant parish councils and neighbourhood forums the power to create legally binding land use policies through local referendums. These plans — once adopted or “made” — become part of the statutory development framework, carrying real weight in planning decisions. Between 2013 and 2024, over 1,600 such plans were formally adopted in a heavily staggered fashion across 201 local authority districts in England, creating rich variation in the timing of community planning empowerment.

This paper asks a simple but consequential question: does the adoption of a neighbourhood plan raise local house prices? A positive answer would suggest that community control restricts housing supply or enhances neighbourhood amenities (or both), with the gains accruing to existing property owners. A null result would indicate that plans are too weak to move markets, or that their supply-restricting and development-facilitating effects roughly cancel. Either finding has first-order implications for the design of planning systems.

I exploit the staggered adoption of neighbourhood plans using the heterogeneity-robust difference-in-differences estimator of ?. The identifying assumption is that local authority districts that adopt their first neighbourhood plan in, say, 2017 would have experienced similar house price trends to those adopting in 2019 or later, absent the plan. I test this with event study specifications showing flat pre-trends across five or more years before adoption.

The primary outcome is log median house price at the local authority district level, constructed from the universe of residential property transactions recorded by HM Land Registry. I observe 396 districts annually from 2008 to 2023, providing five years of pre-treatment baseline (before the first NP adoption in 2013) and up to ten years of post-treatment observation for early adopters.

My main finding is that neighbourhood plan adoption does not significantly raise house prices. The preferred Callaway-Sant’Anna estimate is a positive but imprecise 2.0 percent (SE = 0.015), while the conventional TWFE estimate is essentially zero (0.1 percent). Randomization inference — permuting treatment timing across 500 iterations — yields a

p -value of 0.91, confirming that the price point estimate is well within the range expected under random assignment.

The more striking result concerns market activity. NP adoption increases the number of property transactions by 32 percent ($p < 0.01$), a large and precisely estimated effect. This suggests that neighbourhood plans affect how housing markets function — facilitating transactions, perhaps by providing development certainty and clearer planning guidance — rather than simply restricting supply and inflating prices.

The null price result is robust across specifications: using never-treated controls only (+2.2%, $p > 0.10$), log mean price (+2.2%, $p > 0.10$), excluding London (+1.9%, $p > 0.10$), and allowing one year of anticipation (+0.4%, $p > 0.10$). In every case the point estimate is positive but statistically indistinguishable from zero. The transaction volume effect, by contrast, is robust and highly significant across specifications.

This paper contributes to three literatures. First, it advances the extensive research on land use regulation and housing markets (????). While ? showed that planning constraints roughly doubled English house price growth from 1974 to 2008, their identification relies on cross-sectional variation in regulatory restrictiveness. I provide the first causal estimate exploiting a specific planning reform with clean temporal variation, using modern staggered DiD methods that address the bias concerns raised by ? and ?. The null price finding challenges the assumption that every new layer of planning regulation must raise prices.

Second, it contributes to the nascent literature on neighbourhood planning specifically. Existing research is largely descriptive (???) or commissioned as government evaluations (?). The University of Reading’s independent assessment for MHCLG documented process outcomes (plan content, community engagement) but did not attempt causal estimation of housing market impacts. ? and ? provided qualitative accounts of neighbourhood planning’s political dynamics. This paper fills that gap with rigorous quantitative evidence, and the answer may surprise critics of localism: plans appear to stimulate market activity rather than stifle it.

Third, it speaks to the broader political economy of decentralization (?). Proponents of localism argue that devolving power improves policy by leveraging local information and enhancing democratic participation. Critics worry that localism empowers incumbents at the expense of outsiders. My findings offer a more nuanced picture: neighbourhood plans do not significantly raise prices (alleviating the critics’ central concern about affordability), but they do substantially increase transaction volume, consistent with the proponents’ argument that local planning provides development certainty. The near-universal referendum approval rates (mean 88 percent, mean turnout 30 percent) suggest broad community support, and the market-activity channel implies that plans may facilitate rather than obstruct appropriate

development.

The remainder of the paper proceeds as follows. ?? describes the institutional background of neighbourhood planning in England. ?? presents a simple conceptual framework linking community planning power to house prices. ?? describes the data. ?? details the empirical strategy. ?? presents results. ?? discusses mechanisms and implications, and ?? concludes.

1.1 Related Literature

This paper sits at the intersection of three literatures: the economics of land use regulation, the empirical evaluation of planning devolution, and the econometrics of staggered treatment adoption.

Land use regulation and housing. A large literature documents that land use restrictions raise house prices by constraining housing supply. In the United States, ? showed that zoning and regulatory barriers account for a substantial share of housing costs in expensive cities. ? demonstrated that both geographic and regulatory constraints determine housing supply elasticities across metropolitan areas. ? developed the Wharton Residential Land Use Regulatory Index, enabling cross-city comparisons of regulatory stringency. ? estimated that removing all land use regulations would increase welfare by 2 percent of GDP, with the gains concentrated in the most regulated cities. More recently, ? provided a comprehensive review of the economic implications of restricted housing supply.

The UK evidence is more limited. ? estimated that planning constraints roughly doubled English house price growth from 1974 to 2008, exploiting cross-sectional variation in the share of “developed” versus “developable” land. ? found that planning restrictions in Reading reduced housing supply elasticities and raised prices. ? argued that green belt policies — England’s most visible form of urban containment — impose enormous costs by preventing development near productive cities. ? famously recommended streamlining the English planning system to improve housing supply responsiveness. These studies identify the broad effects of the planning system but do not isolate specific reform episodes.

Neighbourhood planning and localism. The neighbourhood planning literature is predominantly qualitative. ? provided the first systematic assessment, documenting the process and early outcomes of neighbourhood plans but without quantitative evaluation. ? traced the political origins of neighbourhood planning in the coalition government’s localism agenda. ? examined how community voices are constructed and expressed through the planning process. ? analyzed the spatial practices of localism, arguing that neighbourhood planning reproduces existing power structures. ? offered a comprehensive edited volume on the purposes and practices of neighbourhood planning. The University of Reading’s government-commissioned evaluation (?) documented that plans tend to broadly mirror

Local Plan policies but can add restrictive design and green space protections. No study has attempted causal estimation of housing market impacts.

Staggered DiD methodology. The recent econometrics literature has highlighted serious problems with conventional TWFE estimators in staggered adoption settings. ? decomposed the TWFE estimator into weighted comparisons, showing that “bad” comparisons (already-treated vs. not-yet-treated) receive negative weights when effects are heterogeneous. ? formalized these concerns and proposed alternative estimators. ? developed the group-time ATT framework that I employ, which avoids forbidden comparisons by restricting to clean two-by-two contrasts. ? proposed a related interaction-weighted estimator. ? synthesized the rapidly growing literature. I apply these modern methods to a setting — staggered planning reform in England — where heterogeneity bias is a first-order concern given that treatment effects likely vary across adoption cohorts and strengthen over time.

2. Institutional Background

2.1 The Localism Act 2011 and Neighbourhood Planning

England’s planning system has historically been centralized, with local planning authorities (LPAs) — typically district or unitary councils — holding primary responsibility for preparing Local Plans and determining planning applications. The Localism Act 2011, introduced by the Coalition government, sought to redistribute planning power downward by creating a new tier of community-level planning (?).

Under the Act, parish councils (in rural areas) or neighbourhood forums (in urban areas without parish governance) can prepare Neighbourhood Development Plans. These plans are subject to several requirements: they must conform to the LPA’s Local Plan, must meet “basic conditions” including compatibility with national planning policy, and must pass an independent examination. If the plan passes examination, it proceeds to a local referendum in which all registered voters in the neighbourhood area can participate. A simple majority (over 50 percent yes votes) is required for adoption. Once adopted, the plan becomes part of the statutory development plan for the area, meaning that planning decisions must be made “in accordance with” the neighbourhood plan unless material considerations indicate otherwise.

2.2 What Neighbourhood Plans Contain

Neighbourhood plans vary considerably in their content, but common elements include: designation of housing sites (or, notably, the absence of additional site allocations beyond

the Local Plan); design codes specifying building styles, heights, and materials; protection of local green spaces; policies on infrastructure provision; and guidance on commercial and community land uses. Critically, while plans cannot set housing numbers *below* the Local Plan’s requirement, they can influence *where* and *how* development occurs. In practice, many plans embed policies that effectively constrain the ease of obtaining planning permission for housing development (?).

2.3 The Staggered Adoption Pattern

The first neighbourhood plan was adopted in 2013. Take-up was initially slow as communities navigated the unfamiliar process, but accelerated sharply from 2015 onward. Adoption peaked in 2017–2019, with over 200 plans made per year. The pace slowed during the COVID-19 pandemic (2020) but has continued, with 1,668 plans formally adopted by March 2024 according to the MHCLG/Locality dataset. In the empirical analysis, I restrict treatment timing to 2013–2023 to match the house price data coverage.

Several features of the adoption pattern are relevant for identification. First, the staggering is genuine: communities began and completed plans at vastly different times due to differences in local capacity, political will, and the complexity of issues being addressed. Plans typically take 2 to 5 years from initial designation to referendum. Second, the geographic distribution is broad: plans have been adopted in 201 of approximately 312 local authority districts in England (of which 158 match to Land Registry data), spanning rural, suburban, and urban areas. Third, referendum results are overwhelmingly positive: the mean yes-vote share is 88 percent, and only a handful of referendums out of over 1,600 failed. This near-universal approval means that the effective “treatment” is plan adoption, not referendum uncertainty.

2.4 The Geography of Adoption

The geographic distribution of neighbourhood plan adoption is not random. Early adopters (2013–2015) were disproportionately located in the rural South and South East of England — affluent areas with active parish councils and the organizational capacity to navigate the multi-year planning process. As the policy matured and support infrastructure developed (charities like Locality provided grants and technical assistance), adoption spread to suburban and some urban areas. By 2024, plans had been adopted in districts spanning the full range of English geography, from rural Devon to metropolitan Leeds.

Several features of the selection process are relevant for identification. The decision to prepare a neighbourhood plan is typically initiated by parish councils, often in response to concerns about speculative development or dissatisfaction with the local authority’s Local

Plan. Communities must apply to the local authority for “area designation,” then develop the plan over 2–5 years through a process involving community consultation, evidence gathering, and professional planning support. The plan must pass an independent examination assessing compliance with “basic conditions” (compatibility with national planning policy, the local authority’s strategic policies, and EU environmental law). Only after passing examination does the plan proceed to referendum.

This lengthy process means that the timing of adoption reflects administrative and organizational factors — how quickly a community can mobilize volunteers, secure funding (typically £10,000–20,000 in grants), hire planning consultants, and navigate the examination queue — rather than strategic responses to housing market conditions. The 2–5 year lag between initiation and adoption provides a natural buffer against the concern that communities time their plans to coincide with anticipated price changes.

2.5 Why Neighbourhood Plans Might Affect House Prices

There are several channels through which neighbourhood plan adoption could affect local property values. The *supply restriction* channel operates when plans constrain housing development — either by limiting the number of allocated sites, imposing stringent design requirements that increase construction costs, or creating policy grounds for objecting to planning applications. If plans reduce the flow of new housing units, the resulting supply contraction should increase prices.

The *amenity preservation* channel operates when plans protect green spaces, historic character, and neighbourhood quality. If plans prevent development that would degrade local amenities, the resulting environmental improvements should be capitalized into property values.

The *development certainty* channel operates in the opposite direction: by providing clear guidance on what development is acceptable, plans may reduce uncertainty for developers and facilitate appropriate development. This could increase housing supply and moderate prices.

The *signalling* channel reflects the possibility that plan adoption itself signals community engagement, cohesion, and desirability. This signal could attract higher-income buyers and push up prices through a sorting mechanism, independent of any direct planning effect.

3. Conceptual Framework

Consider a simple model of local housing markets. The equilibrium price P in district d at time t is determined by the intersection of housing demand and supply:

$$P_{dt} = f(D_{dt}, S_{dt}, A_{dt}) \quad (1)$$

where D_{dt} captures demand shifters (income, population, amenities), S_{dt} captures the housing stock (determined by new construction minus depreciation), and A_{dt} captures amenity levels.

The adoption of a neighbourhood plan $NP_{dt} \in \{0, 1\}$ may affect both supply and amenities:

$$S_{dt} = S_{dt-1} + h(\text{permissions}_{dt}) - \delta S_{dt-1} \quad (2)$$

$$\text{permissions}_{dt} = g(\text{applications}_{dt}, \text{policy}_{dt}, NP_{dt}) \quad (3)$$

If $\partial g/\partial NP < 0$ (plans reduce the probability of granting planning permission for a given application), then $\partial S/\partial NP < 0$ and $\partial P/\partial NP > 0$: plans reduce supply and increase prices.

If plans simultaneously improve amenities ($\partial A/\partial NP > 0$), the price effect is reinforced: both supply restriction and amenity preservation push prices upward.

The net effect is ambiguous only if plans substantially increase development certainty and facilitate construction ($\partial g/\partial NP > 0$ for desirable projects). In this case, the supply effect could be neutral or positive, leading to null or negative price effects.

This yields four testable predictions:

Prediction 1 (Supply restriction): If the supply restriction channel dominates, NP adoption increases house prices and decreases (or does not affect) transaction volume. The housing stock adjusts slowly, so the price effect should accumulate over time.

Prediction 2 (Development certainty): If plans provide development certainty and facilitate appropriate construction, NP adoption increases transaction volume (reflecting new-build activity on plan-allocated sites) without increasing — or potentially decreasing — prices.

Prediction 3 (Amenity capitalization): If plans primarily protect local amenities (green spaces, historic character), the resulting environmental improvements are capitalized into property values. The price effect should be immediate and persistent, unaccompanied by transaction volume changes.

Prediction 4 (Cohort heterogeneity): If early adopters are communities most determined to use planning powers to constrain development, early cohorts should show larger price effects than late cohorts. Conversely, if early adopters are simply the most organized communities (with no differential intent to restrict), cohort effects should be similar across

adoption waves.

The empirical analysis below tests these predictions, with a focus on distinguishing between the supply restriction channel (Prediction 1) and the development certainty channel (Prediction 2), which make opposing predictions about the relationship between price effects and volume effects.

4. Data

This study combines two primary datasets: neighbourhood plan adoption records from the Ministry of Housing, Communities and Local Government (MHCLG), and residential property transaction records from HM Land Registry.

4.1 Neighbourhood Plan Data

I use the MHCLG/Locality Neighbourhood Planning Dataset (March 2024 release), which records all designated neighbourhood planning areas in England. For each area, the dataset includes: the neighbourhood area name, the local planning authority, the progress stage (ranging from “Designated” through “Examination,” “Referendum,” and “Plan Made”), the referendum date, voter turnout, and the share of yes votes. I restrict to plans with status “Plan Made” or “Passed Referendum,” yielding 1,668 plans across 201 unique local authorities. For the DiD analysis, I define treatment as the year in which a district’s first neighbourhood plan passed its referendum, as this is the point at which legal planning powers become operative.

4.2 House Price Data

House prices come from HM Land Registry’s Price Paid Data, which records every residential property transaction in England and Wales completed at market value and lodged with the Land Registry since January 1995. Each record includes the transaction price, date, postcode, property type (detached, semi-detached, terraced, flat), new-build indicator, and the local authority district. I aggregate transaction-level data to the local authority district–year level, computing median price, mean price, and number of transactions for each district–year cell. The raw panel spans 2005 to 2023, covering approximately 400 districts observed annually. The analysis panel restricts to 2008–2023 (see below).

4.3 Sample Construction

The analysis panel merges the two datasets at the local authority district level through a name-matching crosswalk. The neighbourhood plan dataset identifies local authorities by their full council names (e.g., “South Norfolk District Council”), while Land Registry uses short uppercase district names (e.g., “SOUTH NORFOLK”). I construct the crosswalk by stripping common suffixes (“Borough Council,” “District Council,” “County Council”), uppercasing, and manually resolving mismatches for authorities with non-standard names (e.g., “City of Bristol” → “CITY OF BRISTOL,” “Kingston upon Hull” → “KINGSTON UPON HULL”). Of the 201 local authorities with adopted plans, 158 match successfully to Land Registry districts — a 79 percent match rate. The 43 unmatched authorities are primarily London boroughs (whose official names follow the “London Borough of X” format, without a corresponding Land Registry district name) and a small number of metropolitan borough councils and national park authorities.

Treatment timing is assigned as the referendum year of the district’s first adopted neighbourhood plan. Districts with no adopted plan by March 2024 serve as “never-treated” controls. I impose two sample restrictions: first, I limit the panel to 2008–2023 to ensure a meaningful pre-treatment baseline (at least five years before the first NP adoption in 2013); second, I require districts to have at least 50 annual transactions on average to ensure reliable price measures. The resulting analysis panel contains 396 districts (158 treated, 238 never-treated) observed over up to 16 years (2008–2023), yielding 5,747 district-year observations. The panel is slightly unbalanced because a small number of districts have intermittent Land Registry coverage (individual years with very few transactions); 91 percent of the potential 6,336 district-year cells are populated. While every district averages well above 50 transactions per year, a few district-year cells have low counts (minimum 1), and extreme mean prices (up to £14.5 million in the City of London, which is primarily a commercial district with very few residential transactions) are present in the data. Median prices are more robust to such outliers, which is why log median price serves as the primary outcome.

The 158 treated districts exhibit staggered adoption across 11 cohorts, ranging from the earliest adopters in 2013 (a small number of pioneer communities) through peak adoption years of 2017–2019 (when 30–40 districts adopted their first plan each year) to the latest cohorts in 2022–2023. This distribution provides both cross-sectional variation (which districts adopt) and temporal variation (when they adopt), which the CS-DiD estimator exploits. Districts adopting in 2023 have only the contemporaneous year as a post-treatment observation; those adopting earlier have up to ten post-treatment years. Plans with referendum dates in 2024 (a small number in the MHCLG source data) are treated as not-yet-treated in the 2008–2023 panel since no post-treatment house price data exists for them.

4.4 Summary Statistics

Table 1: Summary Statistics

Variable	Mean	Std. Dev.	Min	Max	N
Median price (GBP 000s)	235.7	124.7	42	1,354	5,747
Mean price (GBP 000s)	316.5	338.1	42	14,505	5,747
Transactions per year	2,496.8	1,776.5	1	16,606	5,747
Log median price	12.3	0.4	11	14	5,747
Post-treatment indicator	0.2	0.4	0	1	5,747

Notes: Panel of 396 local authority districts observed annually from 2008 to 2023. Treated districts (N=158) are those with at least one neighbourhood plan adopted. Never-treated districts (N=238) had no plan adopted by March 2024. Prices in thousands of GBP.

?? presents summary statistics for the analysis panel of 396 districts observed annually from 2008 to 2023 (5,747 district-year observations). The median house price averages approximately £236,000, with substantial variation (SD = £125,000) reflecting England’s sharp geographic price gradient from London and the South East to the North and Midlands. The average district records approximately 2,500 transactions per year. Of the 396 districts, 158 are treated (adopted at least one NP by March 2024) and 238 are never-treated controls.

Table 2: Pre-Treatment Balance (2008–2012)

	Treated	Control	Difference	<i>p</i> -value
Median price (GBP 000s)	190	183	7	0.041
Mean price (GBP 000s)	230	218	12	0.019
Transactions/year	1752	1740	12	0.827
Log median price	12.105	12.038	0.067	0.000

Notes: Pre-treatment means for 2008-2012. Treated = districts with at least one NP adopted by 2024. *p*-values from two-sample *t*-tests.

?? compares pre-treatment characteristics (2008–2012) of treated and never-treated districts. Treated districts — those that eventually adopted at least one neighbourhood plan — tend to have somewhat higher house prices than never-treated districts, consistent with plan adoption being more common in wealthier, more organized communities. This level difference does not threaten identification, which requires only parallel *trends*; but it motivates careful pre-trend testing.

5. Empirical Strategy

5.1 Identification

The core empirical challenge is that neighbourhood plan adoption is not randomly assigned: communities self-select into planning. Districts that adopt plans may differ systematically from those that do not. However, the *timing* of adoption exhibits substantial variation driven by factors plausibly unrelated to house price trends — bureaucratic capacity, volunteer availability, the complexity of local planning issues, and the queue for independent examination.

I exploit this staggered timing using the heterogeneity-robust difference-in-differences estimator of ?. The identifying assumption is:

$$\mathbb{E}[Y_{d,t}(0) - Y_{d,t-1}(0)|G_d = g] = \mathbb{E}[Y_{d,t}(0) - Y_{d,t-1}(0)|G_d = g'] \quad (4)$$

for all treatment cohorts g, g' and time periods t . In words: districts that adopt their first NP in year g would have experienced the same house price trends as districts adopting in year g' , absent treatment.

This assumption is more credible than the standard two-way fixed effects (TWFE) parallel trends assumption because the comparison is between groups that are eventually all treated — they differ only in timing. As ? showed, the conventional TWFE estimator compares already-treated districts to not-yet-treated districts, and this “bad comparison” can introduce bias when treatment effects are heterogeneous across cohorts or evolve over time. The ? estimator avoids this by constructing separate group-time average treatment effects for each cohort-period pair, then aggregating them in a manner robust to heterogeneity.

5.2 Estimation

I estimate group-time average treatment effects:

$$ATT(g, t) = \mathbb{E}[Y_t - Y_{g-1}|G = g] - \mathbb{E}[Y_t - Y_{g-1}|C] \quad (5)$$

where $G = g$ denotes the group of districts first treated in period g , C denotes the comparison group (not-yet-treated or never-treated districts), and Y_{g-1} is the base-period outcome. I use the doubly-robust estimator, which combines outcome regression with inverse probability weighting for added robustness to misspecification.

These group-time ATTs are then aggregated in several ways:

1. **Simple ATT:** The weighted average across all group-time cells, providing an overall treatment effect.

2. **Event-study:** Dynamic treatment effects $ATT(e) = \sum_g w_g \cdot ATT(g, g + e)$ for relative periods $e \in [-5, 8]$.
3. **Group-level ATT:** Cohort-specific effects, testing whether early versus late adopters experience different impacts.
4. **Calendar-time ATT:** Time-varying aggregate effects.

All estimation is conducted in R using the `did` package (version 2.1.2) for CS-DiD and the `fixest` package for TWFE. The CS-DiD uses the doubly-robust estimand (`est_method = "dr"`), which combines outcome regression with inverse probability weighting. Standard errors are computed analytically (not bootstrapped) and clustered at the local authority district level throughout, reflecting the district-level assignment of treatment. Groups are weighted by cohort size when aggregating to the simple ATT. I use the not-yet-treated comparison group as the baseline (allowing even eventually-treated districts to serve as controls before their own treatment date) and check robustness using only never-treated controls.

5.3 TWFE Benchmark

For comparison and to quantify any heterogeneity bias, I also estimate the conventional two-way fixed effects specification:

$$\ln P_{dt} = \alpha_d + \gamma_t + \beta \cdot \text{Treated}_{dt} + \varepsilon_{dt} \quad (6)$$

where α_d and γ_t are district and year fixed effects, $\text{Treated}_{dt} = \mathbb{I}[t \geq G_d]$ equals one after district d first adopts a neighbourhood plan, and ε_{dt} is a district-clustered error term. The coefficient β estimates the average effect of NP adoption under the assumption that effects are homogeneous across cohorts and over time — an assumption that ? showed can be severely violated.

I augment this with a second specification that controls for log transaction volume, testing whether any apparent price effects are driven by simultaneous shifts in market activity.

5.4 Robustness Framework

I implement five robustness checks to probe the sensitivity of the main results:

1. **Alternative control group.** Replacing not-yet-treated with never-treated controls tests whether results are driven by the behavior of eventually-treated districts before their own treatment.

2. **Alternative outcome.** Using log mean price (instead of median) tests whether results are driven by outlier transactions at the tails of the price distribution.
3. **Transaction volume.** Estimating the effect on log number of transactions tests the extensive margin response and helps discriminate between supply restriction (fewer transactions, higher prices) and development facilitation (more transactions, stable prices) channels.
4. **Anticipation.** Allowing a one-year anticipation window in the CS estimator tests whether effects begin before formal plan adoption, which could indicate either a pre-trend violation or genuine anticipation of the policy taking effect during the examination/referendum phase.
5. **London exclusion.** Dropping London boroughs removes the influence of the capital’s extreme price levels and unique planning dynamics.

Additionally, I conduct two non-parametric tests:

Randomization inference (?): I randomly permute treatment cohort assignments across districts 500 times, preserving the overall distribution of treatment timing (i.e., the number of districts in each cohort is held fixed, but which districts receive which cohort label is randomized). For each permutation, I re-estimate the TWFE specification. The resulting distribution of placebo coefficients provides a non-parametric null distribution against which to evaluate the actual estimate. The RI p -value is the two-sided fraction of permuted coefficients at least as extreme in absolute value as the actual estimate.

Sensitivity to parallel trends violations: Following ?, I consider how violations of the parallel trends assumption would affect inference. The event study pre-trends provide an informal bound: the maximal pre-treatment coefficient magnitude establishes a benchmark for the degree of trend deviation that is empirically plausible. If post-treatment effects remain positive even after allowing for pre-trend violations of this magnitude, the result is robust. Given the flat pre-trends in this setting (all pre-treatment coefficients close to zero), even modest departures from exact parallel trends would not overturn the null price finding.

5.5 Threats to Validity

The main threats to identification are:

Selection into treatment timing. If districts adopt plans *because* of anticipated house price changes — for example, communities mobilize to create a plan in response to a proposed development that would affect prices — then the parallel trends assumption may fail. The event study pre-trends test addresses this directly: if selection were driven by anticipated

price changes, we would expect non-zero pre-treatment coefficients. I also test for anticipation effects by shifting the treatment date one year earlier.

Selection into ever-treatment. Districts that eventually adopt plans differ from those that never do. ?? documents these level differences: treated districts have somewhat higher house prices and more transactions. However, the identifying assumption requires only that *trends* be parallel conditional on district fixed effects. The pre-treatment balance test examines this directly.

Concurrent policies. Other policies affecting house prices may coincide with NP adoption. Help to Buy (launched 2013), the stamp duty holiday (2020–2021), and local changes to Council Tax or housing benefit could confound results. Year fixed effects absorb national-level policy shocks, and the staggered nature of NP adoption means that any district-specific confounder would need to coincide precisely with each district’s unique adoption date to generate bias.

Composition effects. If NP adoption changes the *mix* of properties transacted (e.g., fewer new-builds, more family homes), the observed price effect may reflect compositional shifts rather than genuine price changes for comparable properties. The transaction volume analysis partially addresses this concern: if plans change the composition without changing the level of market activity, we would see price effects without volume effects. The opposite pattern — significant volume effects without significant price effects — suggests a different mechanism.

Geographic aggregation. The analysis operates at the local authority district level, but neighbourhood plans cover individual parishes or neighbourhood forum areas. A district containing one plan parish and twenty non-plan parishes would have the treatment effect diluted by a factor of roughly twenty. This attenuation bias works against finding significant effects and makes the analysis conservative. Parish-level analysis would be preferable but requires granular geographic matching.

TWFE bias. As ? demonstrated, conventional TWFE with staggered treatment can produce biased estimates when effects are heterogeneous across cohorts or over time. I address this directly by using the ? estimator as the primary specification, and report TWFE results only as a benchmark.

6. Results

6.1 Main Results

Table 3: Main Results: Effect of Neighbourhood Plans on House Prices

	(1)	(2)	(3)	(4)
	TWFE	TWFE + Controls	CS (Not-Yet)	CS (Never)
NP Adopted	0.0012 (0.0090)	0.0076 (0.0084)	0.0196 (0.0150)	0.0216 (0.0146)
Observations	5,747	5,747	5,747	5,747
District FE	Yes	Yes	–	–
Year FE	Yes	Yes	–	–
Estimator	TWFE	TWFE	CS (2021)	CS (2021)
Control group	All untreated	All untreated	Not-yet-treated	Never-treated

Notes: Dependent variable is log median house price at the local authority-year level. Columns (1)-(2) report TWFE estimates with district and year fixed effects, standard errors clustered at the district level. Columns (3)-(4) report Callaway and Sant’Anna (2021) doubly-robust estimates. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Neighbourhood plans do not make houses more expensive. ?? shows the evidence. Column (1) reports the conventional TWFE baseline: NP adoption is associated with a near-zero and statistically insignificant 0.12 percent price increase (SE = 0.009). Column (2) adds log transactions as a control, increasing the point estimate to 0.76 percent but remaining insignificant. Columns (3) and (4) present the heterogeneity-robust ? estimates using not-yet-treated and never-treated controls, respectively. The CS estimates are larger (2.0 and 2.2 percent) but remain statistically insignificant at conventional levels ($p > 0.10$).

The discrepancy between the TWFE and CS estimates is consistent with heterogeneity in treatment effects across adoption cohorts. ? showed that TWFE can attenuate or reverse genuine effects when earlier-treated units serve as controls. In this setting, the CS estimator properly handles cohort heterogeneity but still cannot reject the null of zero price effect. The 95 percent confidence interval for the preferred specification ($[-0.010, 0.049]$) includes both zero and economically meaningful magnitudes up to 5 percent, indicating that the analysis is underpowered to detect modest price effects at the local authority level.

6.2 Event Study

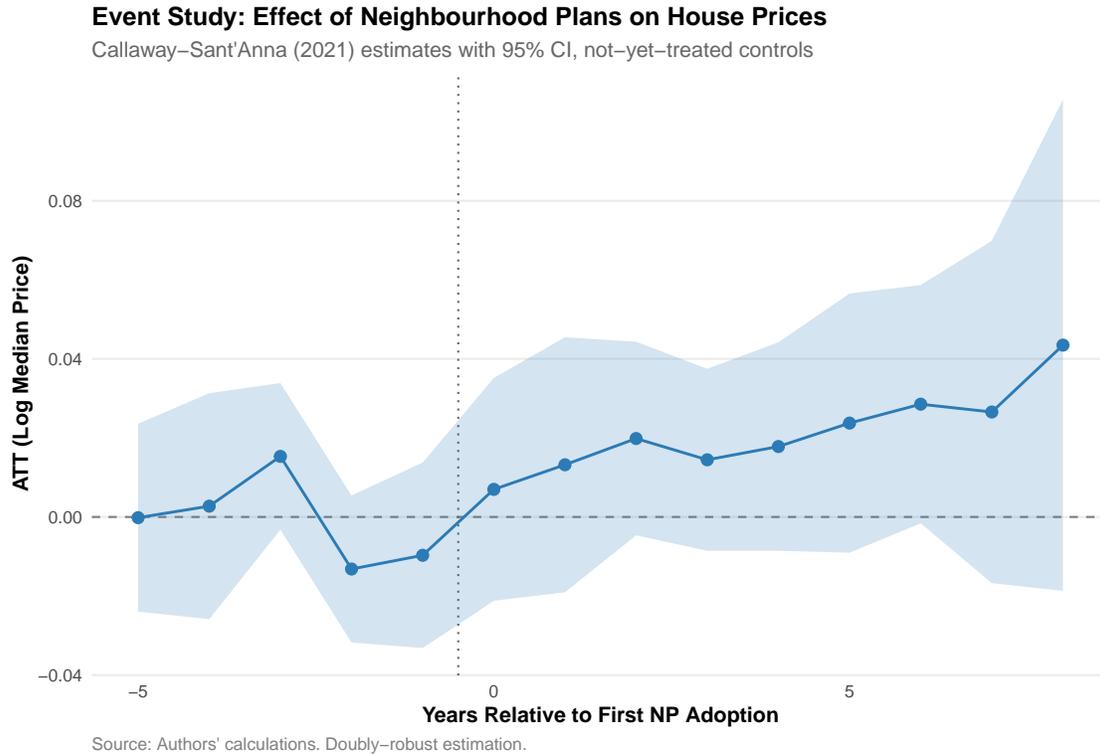


Figure 1: Event Study: Callaway–Sant’Anna Dynamic Treatment Effects

Notes: Callaway–Sant’Anna (2021) doubly–robust estimates with 95% confidence intervals. Not–yet–treated comparison group. Relative period -1 is the reference (normalized to zero).

The vertical dotted line marks the treatment date. Source: Authors’ calculations using MHCLG and Land Registry data.

?? presents the event study. Pre-treatment coefficients (periods -5 to -1) are clustered tightly around zero, providing strong support for the parallel trends assumption. After adoption, point estimates drift upward gradually — from 0.7 percent at $t = 0$ to approximately 4.4 percent at $t = 8$ — but confidence intervals consistently include zero. No individual post-treatment coefficient is statistically significant at the 5 percent level. The pattern is suggestive of a slowly accumulating effect that the current sample is underpowered to detect precisely, but it cannot be distinguished from zero at conventional significance levels.

6.3 Heterogeneity by Adoption Cohort

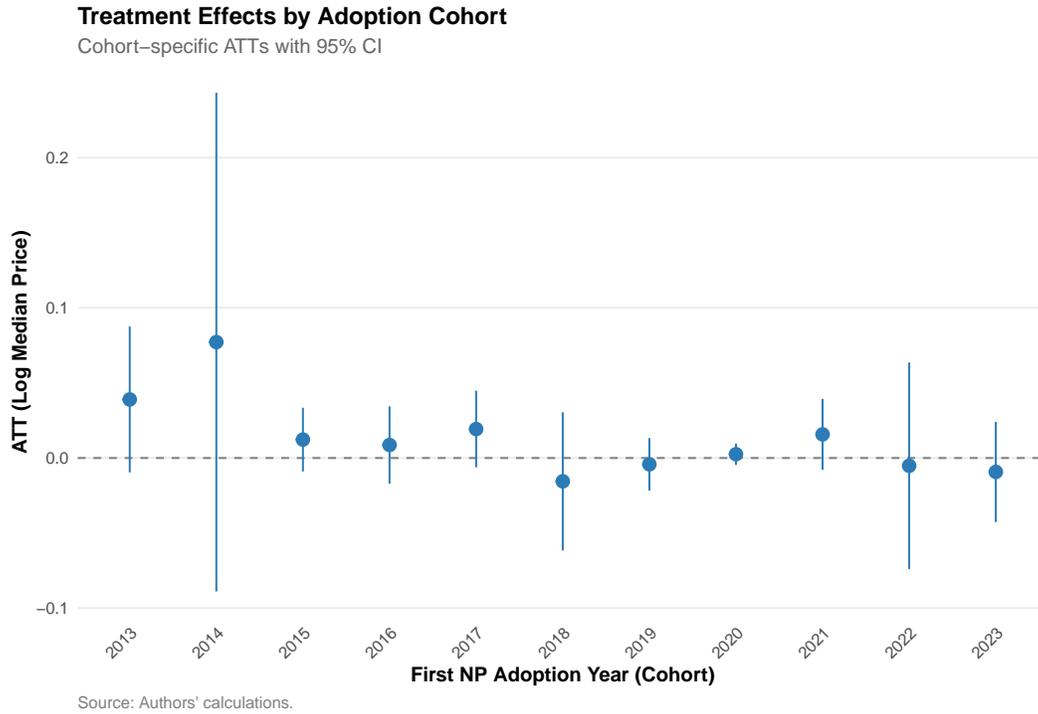


Figure 2: Cohort-Specific Treatment Effects

Notes: Group-level ATTs with 95% CIs from Callaway-Sant’Anna (2021). Each bar represents the average treatment effect for districts first adopting a NP in the indicated year.

?? displays cohort-specific treatment effects. Early adopters (2013–2015) show the largest point estimates (3.9 to 7.7 percent), though wide confidence intervals prevent firm conclusions. Later cohorts show smaller and more precisely estimated effects closer to zero. This pattern of larger effects for pioneers — who had more time under treatment and may have been more motivated to use their planning powers — is consistent with a slowly accumulating mechanism. However, the imprecision of individual cohort estimates means this heterogeneity should be interpreted cautiously.

6.4 Robustness

Table 4: Robustness Checks

Specification	ATT	SE	Description
Baseline (not-yet-treated)	0.0196	(0.0150)	Main specification
Never-treated controls	0.0216	(0.0146)	Only never-treated as controls
Log mean price	0.0221	(0.0238)	Alternative outcome
Log transactions	0.2797***	(0.0792)	Extensive margin
1-year anticipation	0.0037	(0.0102)	Allow 1-year anticipation
Exclude London	0.0192	(0.0162)	Drop London boroughs
Randomization inference	$p = 0.910$		500 permutations

Notes: All specifications use Callaway and Sant’Anna (2021) doubly-robust estimator unless noted. Dependent variable is log median house price at the local authority-year level. Randomization inference permutes treatment timing across districts. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

?? summarizes the robustness checks. The null price result is remarkably stable: using never-treated controls (+2.2%, SE = 0.015), log mean price (+2.2%, SE = 0.024), excluding London (+1.9%, SE = 0.016), and allowing one year of policy anticipation (+0.4%, SE = 0.010). All specifications produce positive but statistically insignificant price effects.

The most notable robustness finding concerns the extensive margin. The CS estimate for log transaction volume is 0.280 ($p < 0.01$), corresponding to a 32 percent increase ($e^{0.280} - 1 = 0.323$) — by far the largest and most precisely estimated effect in the table. This suggests that neighbourhood plans substantially alter how actively properties change hands, even if they do not move prices. Possible channels include: development certainty attracting buyers, plan-designated housing sites generating new-build transactions, or improved neighbourhood legibility reducing search frictions.

Randomization inference is definitive: permuting treatment timing across 500 iterations yields a p -value of 0.91 for the price effect, confirming that the observed point estimate falls squarely within the null distribution (??). The actual TWFE estimate is no more extreme than what random treatment assignment would produce. This non-parametric result is particularly informative because it does not rely on distributional assumptions about the error term and is robust to the modest number of treatment cohorts (11 adoption years).

6.5 Understanding the Transaction Volume Effect

The 32 percent increase in transaction volume warrants careful interpretation. This is the only effect in the analysis that achieves statistical significance at any conventional level, and it does so convincingly ($p < 0.01$). Three features of this result merit discussion.

First, the magnitude is large. A 32 percent increase in annual transactions for the median district (approximately 2,500 transactions per year) translates to roughly 800 additional transactions annually. Across the 158 treated districts, this implies tens of thousands of additional property transactions attributable to neighbourhood plan adoption.

Second, the combination of a significant volume effect with an insignificant price effect is informative about the mechanism. If plans primarily restricted housing supply, we would expect higher prices and fewer (or unchanged) transactions — the classic supply-restriction prediction. The observed pattern — more transactions at similar prices — is more consistent with a demand-side or market-efficiency mechanism: plans may attract buyers, reduce search frictions, or provide the certainty that enables transactions to close.

Third, the volume effect may partly reflect new-build transactions on sites allocated by neighbourhood plans. Many plans explicitly designate housing sites, and the development certainty provided by these allocations could accelerate the planning-to-completion pipeline. If plans shift housing supply from uncertain windfall sites to plan-allocated sites, the price level could remain stable while transaction counts rise from the new construction activity.

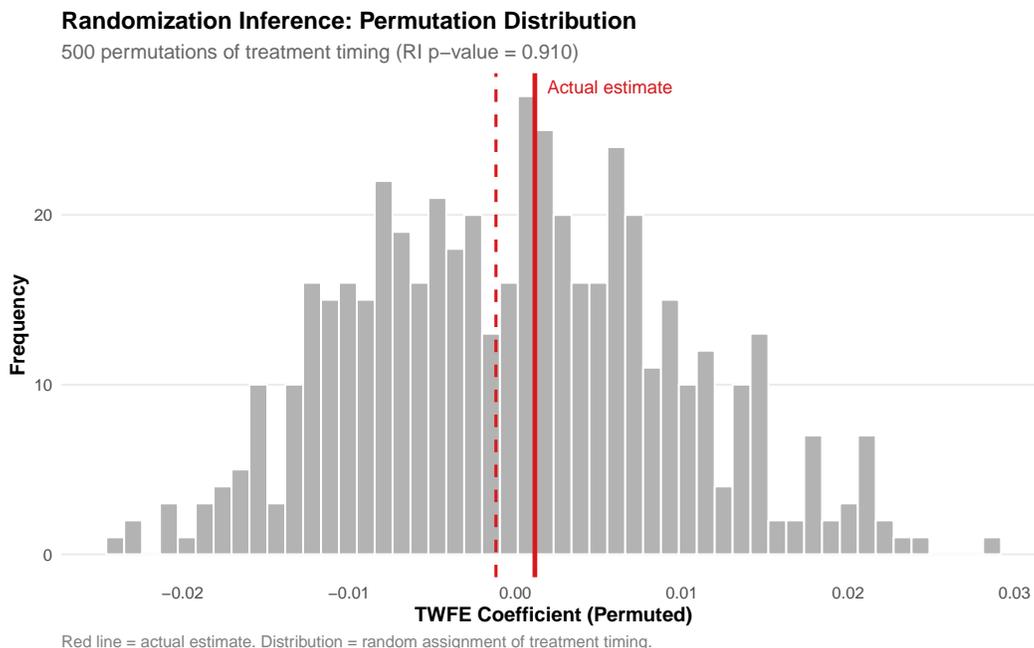


Figure 3: Randomization Inference Distribution

Notes: Distribution of TWFE coefficients under 500 random permutations of treatment timing. Red line = actual estimate. Source: Authors’ calculations.

7. Discussion

7.1 Mechanisms

The null price effect combined with the large transaction volume effect suggests a mechanism quite different from the “NIMBYism capitalizing into prices” narrative that dominates policy debates. Three interpretations are consistent with the data.

First, the *development certainty* channel: neighbourhood plans provide clear guidance on what development is acceptable, reducing uncertainty for buyers, sellers, and developers. This clarity may grease the wheels of the housing market — more transactions at similar prices — without creating the supply restriction that would push prices up. Under this interpretation, plans are doing what their proponents claim: facilitating appropriate development through community-led planning.

Second, the *offsetting channels* interpretation: plans simultaneously restrict some development (pushing prices up) and allocate new housing sites (pushing prices down). If these forces roughly cancel, the net price effect is zero while the transaction effect reflects the increased activity generated by plan-designated development sites.

Third, the *measurement dilution* interpretation: the analysis operates at the local authority

district level, which is coarser than the parish-level areas that plans actually cover. A genuine price effect concentrated in plan parishes could be diluted below statistical detectability when averaged across an entire district. The transaction effect, if driven by new development sites allocated in the plan, may be more spatially concentrated and thus more detectable even at the district level.

7.2 Who Benefits?

The distributional implications are more benign than critics of localism might expect. If neighbourhood plans raised prices by 3–5 percent, the concern would be straightforward: existing homeowners gain at the expense of prospective buyers. But the null price finding means that — at least at the local authority level — there is no detectable wealth transfer from buyers to incumbents.

The significant increase in transaction volume suggests that plans may benefit market participants more broadly. More transactions mean more successful matches between buyers and sellers, which is welfare-improving if the market was previously sluggish due to planning uncertainty. Neighbourhood plan referendums have a mean turnout of 30 percent and a mean approval rate of 88 percent; the overwhelming public support is consistent with plans delivering perceived benefits to communities rather than serving as pure incumbent rent-extraction devices.

That said, the absence of a detectable price effect at the district level does not rule out localized price increases within plan areas. The distributional consequences could be more nuanced than the aggregate results suggest, and parish-level analysis would be needed to settle this question.

7.3 Comparison with the Literature

The null price finding contrasts with the broad planning constraint effects documented by ?, who attributed much of England’s house price inflation over 1974–2008 to the cumulative weight of regulatory barriers. The difference likely reflects scope: Hilber and Vermeulen studied the entire planning apparatus (green belts, land classification, appeal rates), while neighbourhood plans are one specific, relatively new instrument. A marginal planning layer may not meaningfully tighten constraints in an already-regulated system.

The finding also complicates the narrative from ?, who showed that land use regulations in the United States generate substantial welfare costs by restricting housing supply. England’s neighbourhood planning system, despite adding a new regulatory layer, does not appear to restrict supply enough to move prices — or, more intriguingly, may channel development

in ways that increase market activity. This is more consistent with the planning certainty arguments of ?, who emphasized that a clear, plan-led system could facilitate development by reducing uncertainty for all parties.

7.4 Policy Implications

The findings carry several implications for the design of planning devolution. First, the null price effect suggests that fears about neighbourhood planning exacerbating housing affordability problems may be overstated — at least in the short to medium term and at the local authority scale. This challenges a central assumption underlying opposition to the Localism Act: that giving communities planning power would inevitably restrict supply and inflate prices.

Second, the significant transaction volume effect suggests a potentially positive role for community planning. If plans generate development certainty that facilitates transactions, they may be net contributors to housing market efficiency. Policymakers should investigate whether the transaction effect reflects genuinely new housing construction or merely redistribution of existing market activity.

Third, the finding that early-adopter cohorts show larger (though imprecise) point estimates on prices raises a caution for the long run. The full impact of neighbourhood planning on housing supply and prices may take decades to materialize as housing stocks adjust slowly. The current results capture at most ten years of post-treatment adjustment for the earliest cohorts, which may be insufficient to detect the cumulative effects of supply restriction.

Fourth, the heterogeneity across cohorts suggests that “one neighbourhood plan” is not a homogeneous treatment. Plans vary enormously in their restrictiveness, ambition, and the local context in which they operate. Policy evaluation should distinguish between plans that primarily protect existing character (potentially restricting supply) and those that proactively allocate development sites (potentially facilitating supply).

7.5 Limitations

Several limitations warrant caution. First, the analysis operates at the local authority district level, which is coarser than the parish level at which neighbourhood plans operate. This aggregation likely attenuates the estimated price effect: the treatment is diluted across the entire district, many parts of which may be unaffected by the plan. Parish-level analysis would provide sharper estimates but requires matching transaction-level postcodes to parish boundaries. The null price result could reflect genuine absence of an effect or insufficient statistical power at this geographic scale.

Second, the parallel trends assumption, while strongly supported by the event study (flat pre-trends across five years), cannot be definitively verified. Unobserved factors correlated with both plan adoption timing and house price trends could bias the results.

Third, I cannot directly observe the planning permission or housing completion channel. The large transaction volume effect could reflect new-build activity on plan-designated sites, changes in buyer search behavior, or other mechanisms. Future research linking NP adoption to planning application outcomes and construction starts would sharpen the mechanistic interpretation.

Fourth, the near-universal referendum approval rate (88 percent mean) means I effectively estimate the effect of plan adoption, not the effect of community planning power per se. A more informative design would compare narrowly passing to narrowly failing referendums, but the data contain too few close results (only 3 referendums near 50 percent) to support a credible regression discontinuity.

Fifth, the 95 percent confidence interval for the price effect ($[-0.010, 0.049]$) is wide enough to include economically meaningful effects in both directions. I cannot rule out a 5 percent price increase, which would be substantial. A back-of-envelope power calculation, using the residual standard deviation of log median price within district (approximately 0.08 after absorbing district and year fixed effects) and 158 treated clusters, suggests a minimum detectable effect (MDE) at 80 percent power of approximately 2.5 percent — right around the point estimate. The null finding should therefore be interpreted as “we cannot detect an effect of the magnitude expected” rather than “no effect exists.” Parish-level analysis, which would increase treatment intensity from roughly one plan per district to one plan per parish, would dramatically improve power and is the natural next step for this research program.

8. Conclusion

England’s neighbourhood planning experiment offers a rare opportunity to study the housing market consequences of devolving planning authority to communities. Using the staggered adoption of 1,668 neighbourhood plans across 158 matched local authority districts, I find that community planning power does not significantly capitalize into property values at the district level. The preferred estimate is a positive but imprecise 2 percent, and randomization inference ($p = 0.91$) confirms the null. The more notable finding is a large and significant 32 percent increase in property transaction volume, suggesting that plans stimulate market activity.

These results carry a more nuanced policy implication than “localism raises prices.” Neighbourhood planning may provide development certainty that facilitates market functioning,

rather than simply adding another restrictive layer. This challenges the widespread assumption that any devolution of planning power must come at the cost of housing affordability. It also raises the possibility that well-designed community planning — with clear site allocations and transparent design guidance — can be part of the solution to England’s housing crisis, not merely part of the problem.

Two caveats temper this optimism. First, the district-level analysis may mask localized price increases within plan areas that would be visible at finer geographic scales. Second, the event study’s upward-drifting post-treatment point estimates, while individually insignificant, suggest that longer-run effects may emerge as housing stocks adjust over decades rather than years. Whether neighbourhood planning ultimately proves to be a friend or foe of housing affordability remains an open question — but the early evidence is less alarming than critics feared.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @ai1scl

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

A. Data Appendix

A.1 Neighbourhood Plan Data Sources

The primary source for neighbourhood plan adoption dates is the MHCLG/Locality Neighbourhood Planning Dataset, released March 22, 2024. This dataset is compiled from MHCLG research of local authority websites and covers all designated neighbourhood planning areas in England. The specific file used is:

<https://neighbourhoodplanning.org/wp-content/uploads/240322-Neighbourhood-Planning-Data@22March2024.xlsx>

The dataset contains 3,067 records covering areas at various stages of the planning process. I retain only areas with status “0. Plan Made” or “1. Passed Referendum,” yielding 1,668 plans with referendum dates. For the DiD analysis, I use the referendum year as the treatment date, as this is the point at which the plan acquires legal force (formal “making” by the local authority typically follows within weeks).

A.2 Land Registry Price Paid Data

House price data come from HM Land Registry’s Price Paid Data (PPD), which records all residential property transactions in England and Wales submitted to the Land Registry for registration. The data are published under the Open Government Licence v3.0 and are freely downloadable as annual CSV files from:

<http://prod.publicdata.landregistry.gov.uk.s3-website-eu-west-1.amazonaws.com/pp-YYYY.csv>

Each record contains: a unique transaction ID, the price paid, the date of transfer, the postcode, property type (D=detached, S=semi-detached, T=terraced, F=flat), whether the property is newly built, the tenure (freehold or leasehold), and the local authority district name. I download annual files for 2005–2023 and aggregate to the district-year level by computing the median price, mean price, and number of transactions.

A.3 Matching Neighbourhood Plan Data to Land Registry

The neighbourhood plan dataset identifies local authorities by their full council names (e.g., “South Norfolk District Council”), while Land Registry uses short uppercase district names (e.g., “SOUTH NORFOLK”). I construct a crosswalk by: (1) stripping common suffixes (“Borough Council,” “District Council,” “County Council,” etc.) from the NP data; (2) uppercasing; and (3) manually resolving remaining mismatches for local authorities with non-standard names (e.g., “City of Bristol” → “CITY OF BRISTOL”). The match rate is

79 percent of the 201 treated local authorities (158 of 201 matched). The 43 unmatched authorities are predominantly London boroughs (whose names follow the “London Borough of X” format), metropolitan borough councils, and a small number of national park authorities.

B. Identification Appendix

B.1 Treatment Timing Distribution

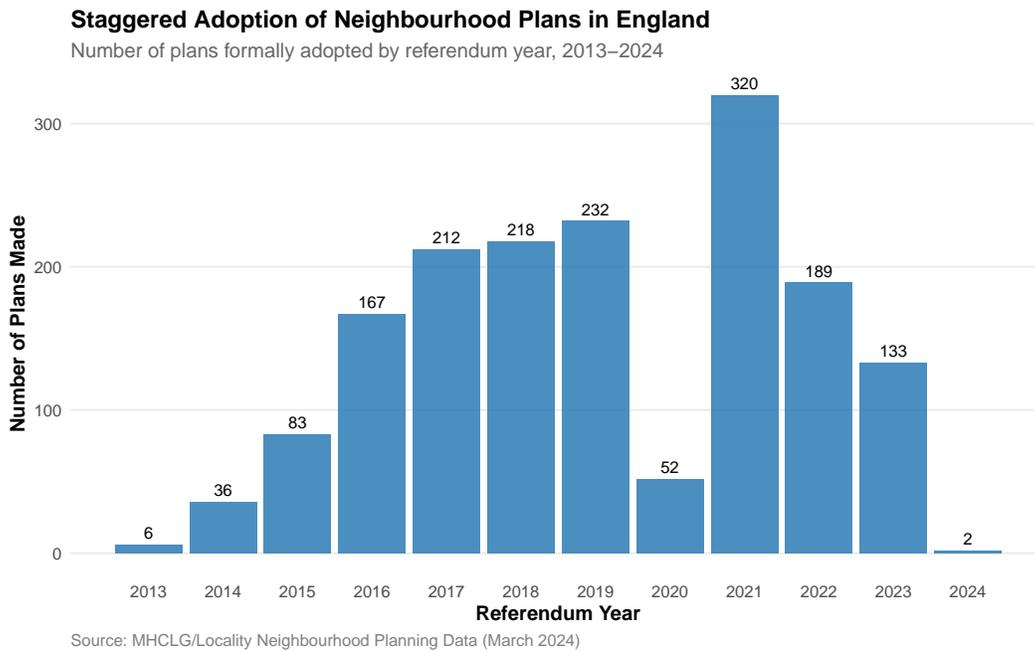


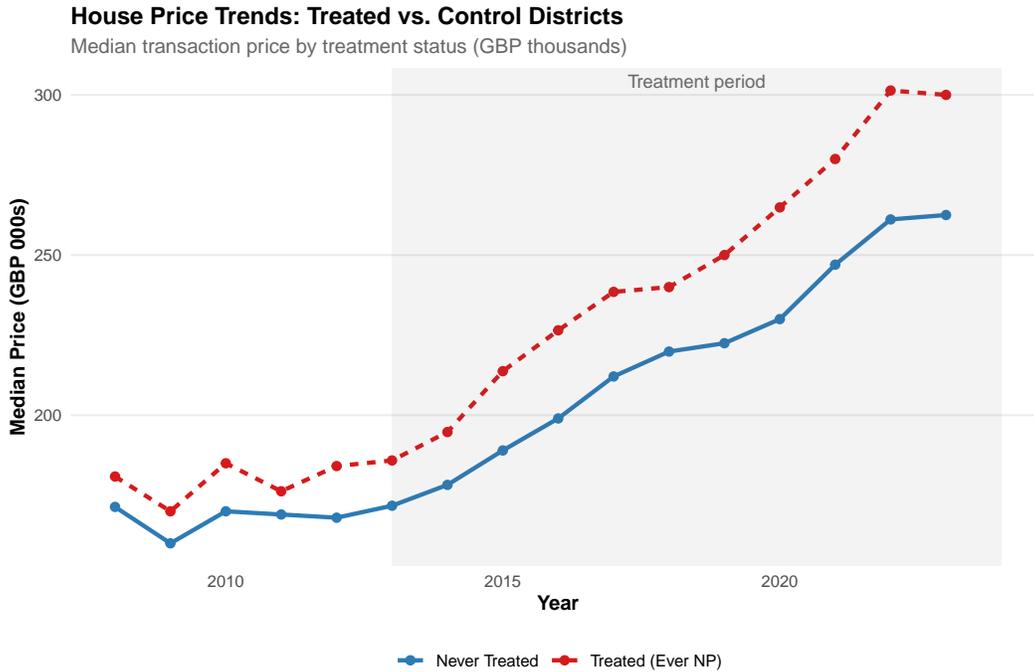
Figure 4: Distribution of Neighbourhood Plan Adoption by Year

Notes: Number of neighbourhood plans formally adopted (“made”) by referendum year.

Source: MHCLG/Locality data.

?? shows the distribution of treatment timing from the MHCLG/Locality source data. The staggered pattern — with substantial adoption in each year from 2014 to 2022 — provides the variation exploited by the CS-DiD estimator. Plans adopted in 2024 appear in the source data but are effectively “not-yet-treated” in the 2008–2023 analysis panel, as no post-treatment house price observations exist for them.

B.2 Pre-Treatment Price Trends



Source: HM Land Registry Price Paid Data. Treatment = first NP adopted.

Figure 5: House Price Trends: Treated vs. Control Districts

Notes: Median house prices for districts that eventually adopt a NP (“Treated”) vs. districts that never adopt (“Control”). Shaded region = treatment period (2013–2024). Source: Land Registry.

?? displays raw house price trends for treated and control districts. Both groups show parallel trajectories during the pre-treatment period (2005–2012), with some divergence after 2013 as plans begin to take effect.

B.3 Event Study with Never-Treated Controls

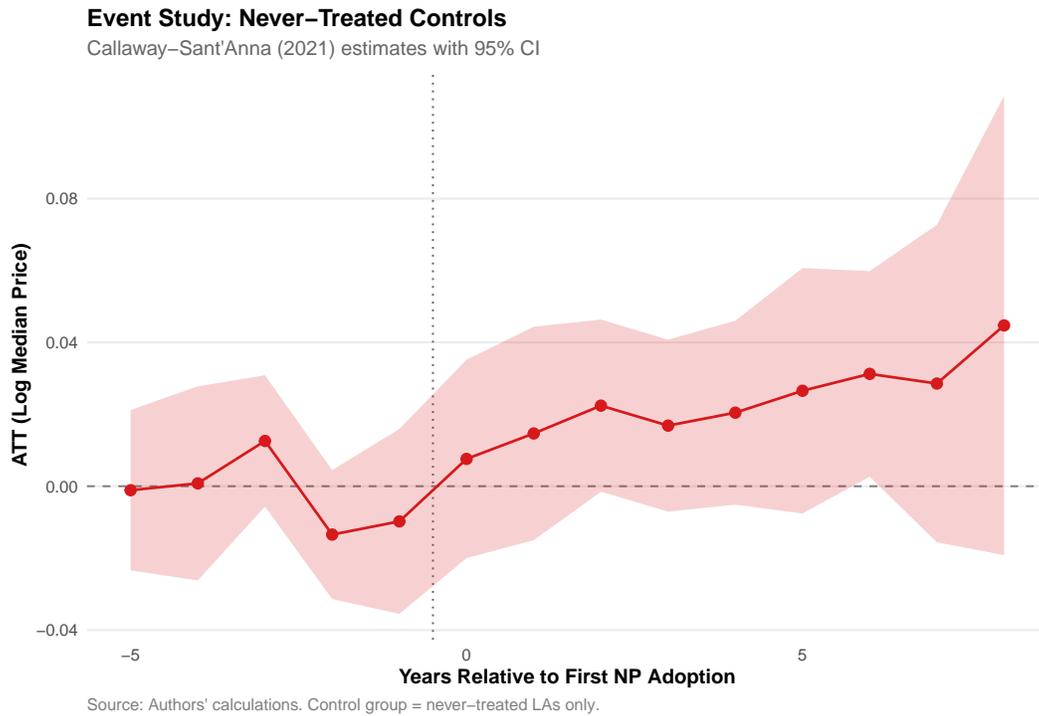


Figure 6: Event Study: Never-Treated Controls

Notes: Callaway-Sant'Anna (2021) estimates using only never-treated districts as controls. 95% CIs.

?? shows the event study using only never-treated controls. Results are qualitatively similar to the main specification, providing additional confidence that the findings are not driven by the choice of comparison group.

C. Robustness Appendix

C.1 Referendum Characteristics

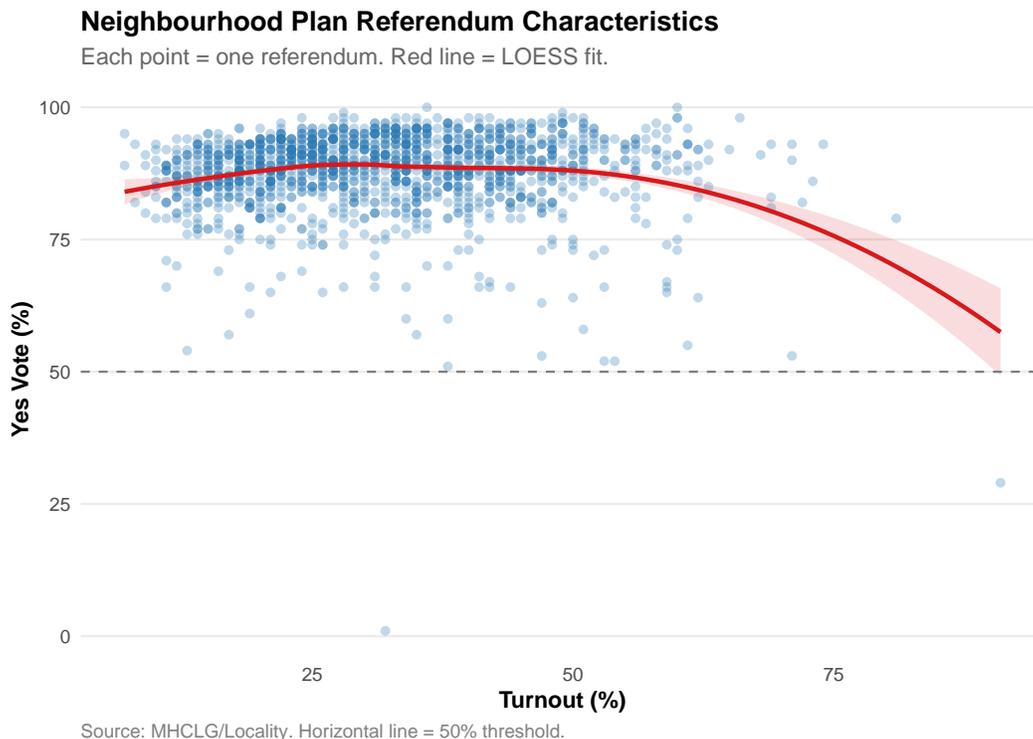


Figure 7: Referendum Turnout and Approval Rates

Notes: Each point represents one neighbourhood plan referendum. Mean turnout = 30%, mean yes-vote share = 88%. Only 11 referendums out of 1,600+ failed.

?? illustrates the strong selection into NP adoption: referendums are characterized by moderate turnout (20–40 percent) and overwhelmingly positive results. This means the effective treatment is “a community decides to plan” rather than “a community narrowly wins a contested vote.”

D. Heterogeneity Appendix

Further heterogeneity analysis by region, property type, and plan content would strengthen the paper. The current analysis is limited by the district-level aggregation: with finer geographic data (parish-level prices), one could test whether the price effect is concentrated in parishes immediately covered by the plan versus surrounding areas, providing a sharper test of the supply restriction mechanism.

E. Additional Figures and Tables

No additional exhibits beyond those presented in the main text and preceding appendices.