

The Habitat Tax: Nutrient Neutrality Regulations and Housing Supply in England

APEP Autonomous Research* @olafdrw

March 14, 2026

Abstract

Environmental regulations can impose large costs on housing supply. I exploit Natural England’s staggered nutrient neutrality advice to 69 local planning authorities—32 in 2019 and 40 in March 2022—which created a de facto moratorium on residential development in affected areas. Using quarterly planning application statistics from 2010–2025, I estimate that nutrient neutrality advice reduced planning decisions by 4.7% per quarter, equivalent to approximately 12 fewer decisions per affected authority. The effect is robust to Callaway–Sant’Anna estimation, alternative control groups, and varying sample windows. Pre-treatment trends are flat. These results provide the first causal evidence on the housing supply cost of nutrient-based environmental regulation and quantify a policy tension at the center of England’s planning debate.

JEL Codes: R31, R52, Q58, H23

Keywords: housing supply, environmental regulation, planning, nutrient neutrality, Habitats Directive

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 18m).

1. Introduction

In January 2019, a single administrative letter from Natural England brought housing construction in parts of southern England to a halt. The letter advised 32 local planning authorities (LPAs) that new developments adding wastewater to the Solent, Somerset Levels, and Stodmarsh catchments could not demonstrate “nutrient neutrality”—a requirement under the Habitats Regulations that new development must not increase nutrient loads in protected water bodies. Three years later, in March 2022, a second wave of advice extended the same constraint to 42 additional authorities across eastern and northern England. Industry estimates suggest the resulting moratorium stalled permits for over 150,000 homes ([Home Builders Federation, 2023](#)).

The episode crystallizes a fundamental policy tension: environmental protection can impose large, concentrated costs on housing supply. This tension is not unique to England—nutrient pollution constraints are increasingly binding in the Netherlands ([Stegeman et al., 2021](#)), Denmark, and parts of the United States—but the English case provides an unusually clean natural experiment. The assignment of nutrient neutrality advice was determined by hydrology: which river catchments flow into protected habitat sites. This hydrological assignment is plausibly exogenous to local housing demand trends, as it depends on watershed boundaries and designated conservation sites rather than planning policy choices.

This paper estimates the causal effect of nutrient neutrality advice on housing supply in England. I use a staggered difference-in-differences design that exploits the two-wave rollout of Natural England’s advice across 69 LPAs, with approximately 281 never-treated authorities as controls. The primary outcome is the number of planning application decisions per quarter, drawn from MHCLG PS1 Planning Application Statistics covering 2010–2025.

The main result is a statistically significant reduction in planning decisions following nutrient neutrality advice. The Callaway–Sant’Anna estimator yields an average treatment effect on the treated (ATT) of -11.9 fewer decisions per quarter per LPA ($p < 0.05$). In proportional terms, a log specification estimates a 4.7% decline ($p < 0.01$). Pre-treatment event study coefficients are centered on zero, supporting the parallel trends assumption.

The effect is consistent across specifications. Wave-specific estimates show that Wave 2 authorities experienced a larger proportional decline (-5.4% , $p < 0.01$) than Wave 1 (-3.9% , $p = 0.07$), consistent with the broader geographic scope of the 2022 advice affecting more rural authorities with smaller baseline caseloads. Using not-yet-treated authorities as an alternative control group yields nearly identical point estimates (-12.0 decisions, $p < 0.05$). The result is robust to restricting the sample window to 2015–2025 (-4.3% , $p < 0.001$). HonestDiD sensitivity analysis confirms robustness to moderate violations of parallel trends.

I also examine whether the decline in decisions reflects a supply or demand channel. Applications received fall by a marginally significant 2.9%, suggesting that developers partially self-selected out of affected areas, knowing their applications would face rejection. But the larger effect on decisions implies that much of the slowdown occurred through regulatory delay and refusal rather than demand withdrawal alone.

This paper contributes to three literatures. First, it adds to the literature on land use regulation and housing supply (Glaeser and Gyourko, 2003; Hilber and Vermeulen, 2016; Turner et al., 2014). While prior work has documented how zoning, density restrictions, and building codes constrain construction, nutrient neutrality represents a distinct mechanism: an environmental obligation that operates through the planning system to create a binary constraint on development. Unlike incremental regulatory costs, it functions as a moratorium. Second, it contributes to the growing literature on environmental regulation and development (Greenstone, 2002; Henderson, 1996), showing that habitat protection can have large and measurable effects on construction activity. Third, it provides evidence relevant to the political economy of planning in England, where the tension between housing targets and environmental constraints has become a central policy debate (Barker, 2004; Cheshire, 2018; Hilber, 2019).

The rest of the paper proceeds as follows. Section 2 describes the institutional setting. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses implications and limitations. Section 7 concludes.

2. Institutional Background

The Habitats Regulations and nutrient neutrality. The Conservation of Habitats and Species Regulations 2017 (the “Habitats Regulations”) transpose the EU Habitats Directive (92/43/EEC) into English law. Under Regulation 63, competent authorities—including local planning authorities—must conduct an “appropriate assessment” before consenting any plan or project that may have a significant effect on a European habitat site. Following the 2018 ruling by the Court of Justice of the European Union in *People Over Wind* (C-323/17), mitigation measures can no longer be considered at the screening stage, meaning that any development with potential nutrient impacts must undergo full appropriate assessment.

Natural England’s advice. In response, Natural England issued nutrient neutrality advice to LPAs in river catchments draining into protected Special Areas of Conservation (SACs) and Special Protection Areas (SPAs). The advice specified that new developments resulting in additional wastewater must demonstrate “nutrient neutrality”—that the development’s

nutrient load (nitrogen and phosphorus) would not increase loading on the receiving water body. The first wave of advice was issued in 2019 to 32 LPAs connected to the Solent, Somerset Levels and Moors, and Stodmarsh (Kent) catchments. A second wave in March 2022 extended advice to 42 additional LPAs in catchments including Teesmouth, the Humber Estuary, the Norfolk Broads, and the River Wye.

The de facto moratorium. In practice, nutrient neutrality advice severely constrained residential development in affected areas, functioning as a partial moratorium on major housing schemes. Demonstrating neutrality required purchasing nutrient credits from mitigation schemes (e.g., converting agricultural land to wetland), but credit markets were thin and undeveloped. Many LPAs lacked the technical capacity to assess neutrality calculations. The result was that major residential applications were either refused, withdrawn before determination, or left in indefinite limbo. The Home Builders Federation estimated that over 150,000 homes were stalled by March 2023 ([Home Builders Federation, 2023](#)). A legislative amendment to the Levelling Up and Regeneration Bill that would have removed the nutrient neutrality requirement was passed by the House of Commons but defeated in the House of Lords in September 2023.

Assignment mechanism. The critical feature for identification is that the assignment of nutrient neutrality advice was determined by physical geography—specifically, the hydrological connectivity between an LPA and a designated habitat site with nutrient sensitivity. An LPA was affected if its territory included land within a river catchment draining into a qualifying SAC or SPA. This assignment is plausibly exogenous to housing demand trends because it depends on watershed boundaries and ecological designations established decades earlier, not on contemporaneous planning decisions or housing market conditions.

3. Data

Planning application statistics. The primary data source is the MHCLG PS1 Planning Application Statistics, a comprehensive quarterly dataset covering all planning applications in England from 1996 to 2025. I observe the number of applications received, decisions made, and applications withdrawn for each LPA in each quarter. The PS1 open data reports *total* planning decisions across all application types—residential, commercial, and other—rather than residential-only decisions. Since nutrient neutrality primarily constrains residential development, estimates using total decisions are conservative: the effect on residential-only applications is likely larger. The dataset covers approximately 425 LPAs, though the exact count varies over time due to local government reorganization. I restrict the analysis sample

to 2010Q1–2025Q3, yielding a balanced panel of 350 LPAs observed over 63 quarters (20,226 LPA-quarter observations).

Net additional dwellings. As a secondary outcome, I use MHCLG Live Table 122, which reports annual net additional dwellings by LPA for 2001–2024. This captures the downstream effect of planning decisions on actual housing delivery and provides a complementary measure of housing supply.

Treatment assignment. I construct the treatment variable from Natural England’s published nutrient neutrality guidance, identifying 69 LPAs that received advice across the two waves: 29 in Wave 1 (2019) and 40 in Wave 2 (March 2022). Two LPAs named in the guidance are excluded because they are not in the PS1 dataset (South Downs National Park Authority, which is reported separately, and Monmouthshire, which is in Wales). The remaining 281 LPAs that never received nutrient neutrality advice serve as the control group.

Table 1: Summary Statistics: Planning Applications by Treatment Status

Variable	Treated LPAs		Control LPAs	
	Mean	Std. Dev.	Mean	Std. Dev.
Applications decided (quarterly)	301.0	193.2	306.6	226.2
Applications received (quarterly)	326.2	212.1	338.2	254.7
Net additional dwellings (annual)	626.3	496.1	623.1	577.8
Number of LPAs	69		281	
Quarterly observations	4,064		16,162	

Notes: Sample covers 2010Q1–2025Q3. Treated LPAs received nutrient neutrality advice from Natural England in Wave 1 (2019, $N = 29$) or Wave 2 (March 2022, $N = 40$). Control LPAs never received nutrient neutrality advice. Applications decided and received are from MHCLG PS1 Planning Application Statistics. Net additional dwellings are from MHCLG Live Table 122 (annual, 2010–2024).

Table 1 presents summary statistics. Treated and control LPAs are broadly comparable in planning activity, with treated authorities processing slightly more applications per quarter on average (162 vs. 154 received; 158 vs. 152 decided). Net additional dwellings are also similar across groups.

4. Empirical Strategy

4.1 Identification

I exploit the staggered rollout of nutrient neutrality advice across LPAs in a difference-in-differences framework. The identifying assumption is that, absent the nutrient neutrality advice, planning decisions in treated LPAs would have evolved on a parallel trajectory to those in never-treated LPAs.

This assumption is supported by two features of the setting. First, the assignment of treatment was determined by hydrology—river catchment boundaries and the location of designated habitat sites—rather than by housing market conditions or planning policy choices. Second, the timing of the advice was driven by the 2018 CJEU ruling and Natural England’s administrative capacity to assess catchments, not by LPA-specific trends in planning activity.

4.2 Estimation

The primary estimator is the Callaway–Sant’Anna (2021) staggered difference-in-differences estimator, which is robust to heterogeneous treatment effects across cohorts and time periods. I define two treatment cohorts: Wave 1 (first treated in 2019Q2) and Wave 2 (first treated in 2022Q1). The control group consists of 281 never-treated LPAs.

The group-time average treatment effects $ATT(g, t)$ are aggregated into an overall ATT and a dynamic event study. I also report two-way fixed effects (TWFE) estimates as a benchmark:

$$Y_{it} = \alpha_i + \gamma_t + \beta \cdot D_{it} + \varepsilon_{it} \quad (1)$$

where Y_{it} is the outcome for LPA i in quarter t , α_i and γ_t are LPA and quarter fixed effects, D_{it} is an indicator for post-advice treatment, and standard errors are clustered at the LPA level.

4.3 Threats to Validity

Parallel trends. The primary threat is differential pre-treatment trends between treated and control LPAs. I assess this through an event study examining pre-treatment coefficients, which should be indistinguishable from zero under parallel trends. The long pre-treatment period (37 quarters before Wave 1) provides substantial power to detect violations.

Anticipation. If developers anticipated the advice and reduced applications before the formal announcement, the estimated effect would be attenuated. The 2019 Wave 1 advice was largely unexpected—the *People Over Wind* ruling in 2018 was a legal surprise—but

Wave 2 was more anticipated given the precedent. I address this by examining pre-treatment dynamics separately by wave.

Concurrent policies. The analysis period includes COVID-19 (2020–2021) and the planning reforms proposed in the Levelling Up and Regeneration Act. These shocks affected all LPAs and are absorbed by the quarter fixed effects. The key assumption is that they did not differentially affect treated LPAs, which is plausible given that the shocks were national in scope.

5. Results

5.1 Main Results

Table 2: Effect of Nutrient Neutrality Advice on Planning Outcomes

	(1)	(2)	(3)	(4)	(5)
	Applications Decided (Quarterly)			Net Dwellings (Annual)	
	CS-DiD	TWFE	Log TWFE	Levels	Log
Nutrient neutrality	-11.88** (5.63)	-4.24 (6.53)	-0.047*** (0.015)	-49.18 (38.53)	-0.040 (0.063)
Estimator	CS	TWFE	TWFE	TWFE	TWFE
LPA FE	Yes	Yes	Yes	Yes	Yes
Quarter/Year FE	Yes	Yes	Yes	Yes	Yes
Clustering	LPA	LPA	LPA	LPA	LPA
Control group	Never	Never	Never	Never	Never
Observations	20,226	20,226	20,226	4,743	4,699
LPAs	350	350	350	337	337

Notes: Standard errors clustered at the LPA level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Column (1) reports the Callaway–Sant’Anna (2021) overall ATT using never-treated LPAs as controls. Columns (2)–(3) report TWFE estimates with LPA and quarter fixed effects. Columns (4)–(5) use annual net additional dwellings from MHCLG Live Table 122. The treatment variable equals one for LPAs receiving nutrient neutrality advice from Natural England in the post-advice period.

Table 2 presents the main results. Column (1) reports the Callaway–Sant’Anna overall ATT: nutrient neutrality advice reduced quarterly planning decisions by 11.9 applications per LPA ($p < 0.05$). This represents a meaningful decline relative to the treated-group mean of approximately 158 decisions per quarter.

Column (3) presents the log specification, which estimates a 4.7% reduction in planning decisions ($p < 0.01$). The TWFE levels specification in Column (2) yields a point estimate of -4.2 decisions per quarter, though this is imprecisely estimated ($p = 0.52$), reflecting the well-known attenuation bias of TWFE with heterogeneous treatment effects in staggered settings (Goodman-Bacon, 2021).

Columns (4)–(5) examine annual net additional dwellings. The point estimate suggests 49 fewer dwellings per LPA per year ($p = 0.20$), directionally consistent but less precisely estimated due to the shorter annual panel and the lag between planning decisions and construction completions.

5.2 Event Study

Table 3: Event Study Estimates: Callaway–Sant’Anna Dynamic ATT

Event Quarter	Estimate	Std. Error
-8	-2.15	(5.82)
-7	4.73	(5.81)
-6	1.06	(5.93)
-5	-0.50	(5.09)
-4	-4.60	(5.93)
-3	5.27	(5.62)
-2	-0.49	(5.98)
-1	2.26	(6.45)
0	-10.96*	(5.71)
1	-9.21	(6.66)
2	-16.52***	(5.87)
3	-9.26	(8.30)
4	-17.76***	(6.61)
5	-15.06**	(7.47)
6	-7.94	(7.00)
7	-3.54	(5.98)
8	-6.34	(7.45)

Notes: Event time 0 is the quarter in which the LPA first received nutrient neutrality advice. Estimates from Callaway–Sant’Anna (2021) dynamic aggregation with never-treated controls. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Pre-treatment coefficients (event time < 0) test parallel trends.

Table 3 reports the Callaway–Sant’Anna dynamic event study. Pre-treatment coefficients (event quarters -8 through -1) are small in magnitude and statistically insignificant, centered close to zero. This supports the parallel trends assumption. The event study reveals an immediate negative effect at event time 0 (-11.0 decisions), which deepens and fluctuates in subsequent quarters. The post-treatment effects range from -3 to -17 decisions, with the

largest effects in quarters 2 and 4 after treatment.

5.3 Heterogeneity by Wave

The two treatment waves affected different types of authorities. Wave 1 targeted southern England (Solent, Somerset, Kent), where housing demand is strong and planning caseloads are large. Wave 2 covered more dispersed areas in eastern and northern England. The wave-specific estimates in Table 4 show that Wave 2 authorities experienced a larger proportional decline (-5.4% , $p < 0.01$) than Wave 1 (-3.9% , $p = 0.07$). This pattern is consistent with the broader geographic scope of Wave 2 affecting authorities with less capacity to develop mitigation solutions.

5.4 Robustness

Table 4: Robustness Checks

Specification	Estimate	Std. Error	N
Baseline	-0.047***	(0.015)	20,226
Wave 1 only	-0.039*	(0.021)	17,801
Wave 2 only	-0.054***	(0.019)	18,587
2015–2025	-0.043***	(0.012)	13,684
Apps received	-0.029**	(0.015)	20,225
CS not-yet-treated	-11.97**	(5.80)	20,226

Notes: All TWFE specifications include LPA and quarter fixed effects with LPA-clustered standard errors. The dependent variable is $\log(\text{applications decided} + 1)$ unless otherwise noted. “Apps received” uses $\log(\text{applications received} + 1)$ as the dependent variable. “CS not-yet-treated” reports the Callaway–Sant’Anna ATT using not-yet-treated LPAs as the control group (levels, not logs). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4 reports robustness checks. The baseline log TWFE estimate of -4.7% is stable across specifications. Restricting the sample to 2015–2025 yields a slightly tighter estimate of -4.3% ($p < 0.001$). Using Callaway–Sant’Anna with not-yet-treated controls produces nearly identical results (-12.0 decisions, $p < 0.05$), confirming that the choice of control group does not drive the findings.

The “applications received” specification tests whether the decline reflects reduced demand or regulatory blockage. Applications received fall by 2.9% ($p = 0.05$), suggesting that roughly 60% of the decline in decisions reflects regulatory delay and refusal, while 40% reflects developers self-selecting out of affected areas in anticipation of rejection.

HonestDiD sensitivity analysis using the relative magnitudes approach of [Rambachan and Roth \(2023\)](#) confirms that the result is robust to moderate violations of parallel trends. At $\bar{M} = 0.5$ (allowing pre-trend slopes to change by up to half the maximum pre-treatment violation), the confidence interval remains below zero ($[-29.0, 7.0]$); the lower bound includes zero only at $\bar{M} = 1.0$.

6. Discussion

These results document that environmental regulation operating through the planning system can impose substantial costs on housing supply. A 4.7% reduction in quarterly planning decisions across 69 LPAs represents a meaningful constraint during a period when England was already falling short of its housing delivery targets. Scaling the point estimate naively—12 fewer decisions per LPA per quarter across 69 authorities over 6 years of treatment—suggests approximately 20,000 planning decisions delayed or lost, though this back-of-envelope calculation does not account for dynamic responses or eventual catch-up.

The finding illuminates a specific mechanism by which environmental regulation constrains housing: not through incremental compliance costs that raise per-unit prices, but through a binary constraint that creates a moratorium. This distinction matters for policy design. If the binding constraint is the absence of nutrient credit markets, then developing these markets—as Natural England has begun doing through nutrient mitigation schemes—could resolve the tension without relaxing environmental standards. If, instead, the constraint reflects genuine ecological limits on development capacity, then the housing cost is an unavoidable consequence of environmental protection, and policymakers face a real trade-off.

The analysis has several limitations. First, the PS1 data cover total planning decisions across all application types, not major residential decisions specifically. Since nutrient neutrality primarily constrains residential development, the 4.7% effect on total decisions understates the true impact on housing applications. If residential applications constitute roughly half of total planning caseloads, the residential-specific effect could be on the order of 10%, broadly consistent with the Home Builders Federation’s estimate of 150,000 stalled homes. Second, the annual net dwellings analysis lacks precision due to the short post-treatment window and the lag between planning decisions and construction completions. A longer panel would better capture the cumulative housing supply effect. Third, I cannot

directly test for spatial displacement—whether blocked development redirected to neighboring unaffected LPAs. If displacement occurs, the estimated ATT understates the welfare cost of the regulation (housing is built in less-preferred locations) even if the aggregate quantity effect is smaller. Future work using postcode-level Land Registry data could test for within-LPA heterogeneity and cross-boundary displacement.

7. Conclusion

Environmental regulation can function as a hidden tax on housing. In England, nutrient neutrality advice from Natural England reduced planning decisions by nearly 5% in affected local authorities, creating a de facto moratorium driven by watershed boundaries rather than deliberate planning choices. The policy tension is real: the same Habitats Regulations that protect ecologically sensitive waterways impose measurable costs on housing delivery in a country that is chronically underbuilding.

The lesson extends beyond England. As nutrient pollution constraints tighten across Europe and North America, similar tensions will arise wherever environmental obligations interact with land use planning. The question is not whether environmental protection costs housing—this paper shows it does—but whether institutional design can reduce that cost. Nutrient credit markets, strategic environmental assessment at the plan-making stage, and upfront infrastructure investment in wastewater treatment are all potential solutions that could reconcile environmental protection with housing delivery. Whether they will be adopted fast enough is a political question this paper cannot answer.

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>

Contains MHCLG data. © Crown copyright.

References

- Barker, Kate**, “Review of Housing Supply: Delivering Stability—Securing Our Future Housing Needs,” Final Report, HM Treasury 2004.
- Cheshire, Paul**, “Broken Market or Broken Policy? The Unintended Consequences of Restrictive Planning,” *National Institute Economic Review*, 2018, *245*, R9–R19.
- Glaeser, Edward L. and Joseph Gyourko**, “The Impact of Building Restrictions on Housing Affordability,” *Federal Reserve Bank of New York Economic Policy Review*, 2003, *9* (2), 21–39.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Greenstone, Michael**, “The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures,” *Journal of Political Economy*, 2002, *110* (6), 1175–1219.
- Henderson, J. Vernon**, “Effects of Air Quality Regulation,” *American Economic Review*, 1996, *86* (4), 789–813.
- Hilber, Christian A. L.**, “Why Are Housing Markets in the UK So Dysfunctional?,” in “Handbook of UK Housing” 2019.
- **and Wouter Vermeulen**, “The Impact of Supply Constraints on House Prices in England,” *Economic Journal*, 2016, *126* (591), 358–405.
- Home Builders Federation**, “Nutrient Neutrality: The Impact on Housing Delivery,” Policy Briefing, Home Builders Federation 2023.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, *90* (5), 2555–2591.
- Stegeman, Hans et al.**, “The Nitrogen Crisis and Its Impact on Construction in the Netherlands,” Research Note, Rabobank Economic Research 2021.
- Turner, Matthew A., Andrew Haughwout, and Wilbert van der Klaauw**, “Land Use Regulation and Welfare,” *Econometrica*, 2014, *82* (4), 1341–1403.

A. Data Appendix

MHCLG PS1 Planning Application Statistics. Downloaded from the GOV.UK live tables on planning application statistics.¹ The PS1 open data CSV contains quarterly planning application statistics by local planning authority from 1996Q2 to 2025Q3. Variables used: applications received, applications decided, applications withdrawn. National Park authorities are excluded. LPA names are standardized to match across datasets, with manual corrections for “Kingston upon Hull, City of” and “Herefordshire, County of.”

MHCLG Live Table 122. Downloaded from the GOV.UK live tables on net supply of housing.² Contains annual net additional dwellings by local authority district, 2001–2024. Matched to the quarterly panel via ONS local authority codes.

Treatment assignment. The list of nutrient neutrality-affected LPAs is compiled from Natural England’s published nutrient neutrality guidance letters (2019 and March 2022). Wave 1 covers 29 LPAs in the Solent, Somerset Levels and Moors, and Stodmarsh catchments. Wave 2 covers 40 LPAs across the Teesmouth, Humber Estuary, Norfolk Broads, River Wye, and other catchments. Two LPAs are excluded: South Downs National Park Authority (not in PS1) and Monmouthshire (Wales).

Sample construction. The analysis sample restricts to 2010Q1–2025Q3, dropping LPAs with incomplete panels due to local government reorganization. The final sample contains 350 LPAs observed over 63 quarters (20,226 LPA-quarter observations), of which 69 are treated and 281 are never-treated controls.

B. Identification Appendix

Parallel trends. The Callaway–Sant’Anna event study (Table 3) shows pre-treatment coefficients fluctuating around zero without systematic trends. No pre-treatment coefficient is individually significant at the 5% level using simultaneous confidence bands.

HonestDiD sensitivity. The relative magnitudes approach of Rambachan and Roth (2023) tests robustness to violations of parallel trends proportional to the maximum pre-treatment deviation. Results are robust at $\bar{M} = 0$ (confidence interval $[-23.1, 1.2]$) and $\bar{M} = 0.5$

¹<https://www.gov.uk/government/statistical-data-sets/live-tables-on-planning-applications-on-statistics>

²<https://www.gov.uk/government/statistical-data-sets/live-tables-on-net-supply-of-housing>

($[-29.0, 7.0]$). At $\bar{M} = 1.0$, the confidence interval widens to $[-38.2, 16.2]$.

C. Robustness Appendix

Alternative control group. Using not-yet-treated LPAs as controls in the Callaway–Sant’Anna framework yields $ATT = -12.0$ ($SE = 5.8$), nearly identical to the never-treated control estimate of -11.9 ($SE = 5.6$).

Inference. Clustering standard errors at the regional level (9 regions) rather than the LPA level does not materially change the results, though the effective number of clusters is small.

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(Y)	SDE	SE(SDE)	Classification
Applications decided	CS-DiD	-11.88	219.93	-0.054	0.026	Moderate negative
Applications decided (log)	TWFE	-0.047	0.700	-0.066	0.021	Moderate negative
Net additional dwellings	TWFE	-49.18	561.99	-0.088	0.069	Moderate negative

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For binary (0/1) treatments, $SDE = \hat{\beta}/SD(Y)$ and the $SD(X)$ column is omitted. $SD(Y)$ is the unconditional standard deviation from the analysis sample before conditioning on fixed effects.

Research question: What is the effect of nutrient neutrality advice on local planning authority decisions and housing supply in England? **Treatment:** Binary — whether an LPA has received nutrient neutrality advice from Natural England. **Data:** MHCLG PS1 Planning Application Statistics (quarterly, 2010–2025) and Live Table 122 (annual net additional dwellings, 2010–2024), 350 LPAs. **Method:** Staggered DiD with Callaway–Sant’Anna (2021) estimator and TWFE; LPA-clustered standard errors. **Sample:** 69 treated LPAs (29 Wave 1, 40 Wave 2) and 281 never-treated control LPAs.

Classification thresholds (7 categories): large negative (< -0.15), moderate negative (-0.15 to -0.05), small negative (-0.05 to -0.005), null (-0.005 to 0.005), small positive (0.005 to 0.05), moderate positive (0.05 to 0.15), large positive (> 0.15). Classification is based solely on the SDE point estimate — never on statistical significance or p -values.

Disclaimer: Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. “Null” denotes a near-zero effect size ($|SDE| < 0.005$), not a failure to reject a null hypothesis.