Regional Science and Urban Economics 121 (2026) 104264

Contents lists available at ScienceDirect

Regional Science and Urban Economics

journal homepage: www.elsevier.com/locate/regec

Sponge city and impermeable housing value: Evidence from London ☆

Yihao Wu

Harvard University, Cambridge, 02138, Massachusetts, USA

ARTICLE INFO

Keywords:
Urban flooding
Green stormwater infrastructure
Property values
Housing market
Difference-in-differences
Sustainable Drainage Systems

ABSTRACT

Climate change has intensified flash flooding, prompting cities worldwide to implement green stormwater infrastructure as a scalable adaptation strategy. This paper evaluates the economic consequences of London's Sustainable Drainage Systems (SuDS) program on the local housing market. Using a staggered difference-in-differences design and an LSOA-by-year panel dataset from 2010 to 2024, I find that neighborhoods receiving green infrastructure record 3.2–3.7% increases in housing prices on average but no discernible rise in transaction volumes. Event study estimates reveal a hump-shaped temporal profile, with effects peaking around the fifth year before partially attenuating. The price responses exhibit a non-monotonic pattern across flood risk categories, largely because infrastructure salience drives value creation more than hydraulic performance. My findings suggest that successful climate adaptation requires balancing engineering effectiveness with public perception, and that green infrastructure can generate sufficient returns to self-finance through property tax revenue.

1. Introduction

Urban flash flooding represents a first-order economic challenge for metropolitan areas worldwide. The interaction between climate-driven precipitation intensification and expanding impervious surface coverage generates substantial welfare losses that extend beyond immediate property and infrastructure damage (Balaian et al., 2024; Kaiser and Akter, 2025). Compared with storm-induced coastal flooding, flash flooding harms more people and inflicts lasting damage in urban areas. Although headline figures associated with urban flooding are typically reported as total direct losses at the city-wide scale, the economic fallout and persistent repercussions are felt quite locally (Costa et al., 2024; Kaźmierczak and Cavan, 2011). Beyond direct damage to property and infrastructure, flash flooding can disrupt local economies through multiple channels from business interruption to supply chain shocks (Jongman, 2018). Furthermore, declining property values due to flooding erode the municipal tax base required for public goods provision, which can trigger a climate-debt doom loop (Zenios, 2024).

Over the past decade, governments have increasingly recognized the potential of natural systems for flood control. Green stormwater infrastructure (GSI) has emerged as a distributed adaptation strategy that mimics natural hydrological processes to manage urban runoff at the source (Finewood et al., 2019). Unlike traditional gray infrastructure that channels stormwater through centralized systems, GSI installations — including rain gardens, bioswales, permeable pavements, and green roofs — absorb and infiltrate precipitation where it falls (Pons et al., 2023). This decentralized approach offers potential cost advantages: individual installations require lower capital investment than sewer expansion, implementation can proceed incrementally aligned with maintenance cycles, and projects generate co-benefits including urban cooling and aesthetic improvements (Chui et al., 2016). However, achieving city-wide flood resilience through GSI requires substantial aggregate investment and ongoing maintenance expenditures, raising questions about optimal financing mechanisms and the distribution of costs and benefits (Cousins and Hill, 2021).

☆ Helpful and much appreciated suggestions, critiques, and encouragement were provided by Rachel Meltzer, Richard Peiser, Siqi Zheng, Edward Glaeser, Avis Devine, and Dongxiao Niu. I am immensely grateful to the Editor and two anonymous referees, whose comments substantially improved the paper. I also thank Betsy Donald, Meric Gertler, Marko Irisarri, Carmen Cantuarias-Villessuzanne, and other participants at the Cambridge Journal of Regions, Economy and Society 2025 Conference and the 15th European Meeting of the Urban Economics Association (UEA) for valuable feedback. I thank Fengyuan Han for excellent research assistance with the geospatial analysis and data processing, and the Greater London Authority and the Environment Agency for providing access to the SuDS retrofit mapping data and flood risk modeling outputs. I gratefully acknowledge financial support from the Salata Institute for Climate and Sustainability and the Harvard Joint Center for Housing Studies. Computational resources were provided by the Harvard Research Computing facility. All errors remain my own.

E-mail address: yihao_wu@gsd.harvard.edu.

https://doi.org/10.1016/j.regsciurbeco.2026.104264
Received 4 August 2025; Received in revised form 1 July 2026; Accepted 13 July 2026
Available online 20 July 2026

The economic evaluation of GSI investments has become increasingly urgent as climate projections indicate continued intensification of extreme precipitation events. Despite widespread implementation across developed cities, particularly in Europe and North America,[1] empirical evidence on GSI's local economic returns remains limited (see Appendix A provides a detailed literature review on related works). This raises a policy-relevant question about the economic benefits of those micro-scale greenery infrastructure: whether it is a low-cost, scalable model with sizeable payoffs, and how these payoffs translate into the local economy? Understanding the capitalization of flash flood risk reduction into property values provides a revealed preference measure of residents' willingness to pay for climate adaptation and the effectiveness of public provision of GSI programs in building urban resilience.

Although recent years have seen a growing academic interest in estimating causal effects of place-based climate adaptation, the primary concern for identification is *selection-on-observable*, particularly in smaller geographic areas with merely a handful of adaptations. Most past studies (see Sohn et al., 2020, Davis et al., 2023 and Park and Kweon, 2025) have used hedonic price models (HPM), estimated using Ordinary Least Squares (OLS), to control for differences in basic property characteristics and estimate the treatment effect of green infrastructure on local housing prices. However, studies in climate justice (Anguelovski et al., 2019; Zuniga-Teran et al., 2021) have noted that the selection and implementation of adaptations are based in part on the neighborhood characteristics (e.g. socioeconomic status, political influence) that are unobserved and not fully captured by analysts, and therefore are not held constant in the estimated HPM. Typically, if high-amenity neighborhoods systematically receive more green infrastructure investment, ordinary least squares estimates will conflate treatment effects with pre-existing differences, biasing estimates upward.

While perfect random assignment of GSI is hard to engineer, addressing this selection bias requires strategies that generate plausibly exogenous variation in treatment assignment. One recent study by Cohen et al. (2025) has posed a clever quasi-experimental method to solve this problem by exploiting the pipeline of planned but not-yet-completed GSI projects as a control group. Their difference-in-differences estimates indicate price premiums of approximately 10 percent in Philadelphia, with effects decaying sharply at greater distances. Several questions with crucial policy and finance relevance remain open notwithstanding. First, the mechanisms through which green infrastructure generates value are still unknown and inconclusive. Second, few studies have examined whether GSI capitalization varies with baseline flood exposure, leaving it unclear whether investments should be targeted to high-risk areas. Third, the property-level treatment definition, while precise, may not correspond to the neighborhood-scale at which infrastructure investment decisions are made and public finance calculations are conducted.

Motivated by these ambiguities, I investigate the capitalization of GSI investments in the UK housing markets. London's gradually enacted Sustainable Drainage Systems (SuDS) program across 33 local authorities induces quasi-experimental variation that comes close to an ideal experiment. Following severe flooding in 2007 and 2011, the Greater London Authority (GLA) initiated a decentralized implementation strategy across 33 local authorities, with installations proceeding according to local capacity, funding availability, and opportunities for integration with scheduled maintenance. Although not random, the staggered rollout, largely shaped by planning pipelines and co-scheduling with programmed works, yields useful quasi-experimental timing variation across broadly comparable areas. Between 2012 and 2024, the program has delivered over 2800 installations with public investment exceeding £20 million. While the benefits of public spending on green infrastructure are taken for granted, its actual effects and economic consequences have not been systematically evaluated. The SuDs program's sporadic, geographically dispersed implementation pattern, combined with the highly localized nature of flash flooding impacts, creates comparability across neighborhoods that received SuDS treatment and those without such an adaptation.

By combining detailed housing transaction records with granular geospatial risk information, I first present evidence of price discounts associated with flash flooding risk in London's housing market through a cross-sectional hedonic price analysis. Properties in high flash flood risk areas (greater than a 3.3 percent annual probability) trade at discounts of 15.0 percent in 2010 and 12.1 percent in 2024, controlling for property characteristics and borough fixed effects. The reduction in flood risk capitalization between periods suggests potential adaptation effects. My main identification strategy then exploits the staggered adoption of the SuDS program through a difference-in-differences design applied to a neighborhood-year panel spanning 2010–2024. Using Layer Super Output Areas (LSOAs) as the unit of analysis, I define treatment as the presence of any SuDS delivery inside an LSOA, with treatment timing determined by the implementation date. The identifying assumption that, absent treatment, property values in treated and control neighborhoods would have followed parallel trends is supported by event study estimates showing no differential pre-trends in the five years preceding adaptation.

I document that GSI investment in London generates measurable economic benefits through property value appreciation even in the absence of major flooding events. The staggered implementation of SuDS across London from 2012–2024 increased housing prices by 3.2–3.7 percent on average, depending on the specification, albeit with no measurable increase in transaction volumes. Event study estimates reveal a distinctive temporal profile: an immediate price response upon installation, partial attenuation by year three, a resurgence to peak effects of approximately 8 percent around year five, and gradual decay thereafter. This hump-shaped trajectory is in accordance with models of belief updating in which households revise expectations about infrastructure durability and flood protection as they observe system performance through subsequent storm events. My results suggest three key insights for scaling green infrastructure investment. First, the substantial property value premiums justify aggressive public investment in neighborhoods with either high flood risk or high visibility, as both generate comparable returns through different channels. Second, the absence of spatial spillovers simplifies value capture mechanisms: tax increment financing or special assessment districts can recover costs from direct beneficiaries without complex multi-jurisdictional coordination. Third, the importance of visibility suggests that showcase projects in prominent locations may catalyze broader public support more effectively than technically optimal but invisible systems.

The rest of the article is organized as follows. The next section describes London's SuDS program, the institutional setting that governs site selection and financing, and the data used to construct neighborhood-level measures of treatment, flood risk, and housing-market outcomes. Section 3 outlines the empirical framework and identification strategy, including the baseline difference-in-differences design, event study specifications, and robustness checks using matching and weighting. Section 4 presents the main results on housing prices and transaction volumes and examines heterogeneity by baseline flood risk. Section 5 investigates mechanisms, focusing on realized hydraulic performance during the July 2021 events, the salience and information channel across SuDS typologies, and the role of greening and spatial spillovers. Section 6 discusses the public-finance and cost–benefit implications of the estimated capitalization effects. Section 7 concludes.

[1] Major green stormwater infrastructure programs in the U.S. include: Philadelphia's *Green City, Clean Waters*, New York City's *NYC Green Infrastructure*, Seattle's *SEA Street* and *RainWise* programs, Chicago's *Green Stormwater Infrastructure Strategy*, and Washington D.C.'s *District Stormwater LLC* and Vancouver's *Rain City Strategy*; In Europe, Copenhagen's *Cloudburst Program*, London's *Sustainable Drainage Action Plan*, and Malmö's *Blue-Green Fingerprints initiative*, etc.

2. Background and data

2.1. Overview of London's SuDS program

London's drainage infrastructure has faced mounting pressure from demographic growth, urban densification, and climate change in recent years (Babovic and Mijic, 2019). The city's combined sewer system, constructed during the Victorian era under Joseph Bazalgette's direction (1864–1877), served a substantially smaller urban population with greater pervious surface coverage. Contemporary urbanization in metropolitan London has fundamentally altered the hydrological characteristics of the area, creating conditions that exceed the design capacity of existing sewage infrastructure (Campos and Darch, 2015). The marginal cost of traditional gray infrastructure expansion to accommodate increased surface water volumes has become prohibitive. Thus, the recourse to nature-based solutions as a means of managing stormwater has become inevitable.

In 2010, with public funding from the Department for Environment, Food and Rural Affairs, London established the *Drain London Partnership* to deal with surface water flood risk and provide sustainable drainage retrofitting through Sustainable Drainage Systems (SuDS). Several SuDS projects were piloted thereafter and have been delivered gradually since 2012, while the program gained significant momentum in 2015 when the London Sustainable Drainage Action Plan (LSDAP) was officially published (Fig. 1). According to an announcement by the then Mayor of London, Boris Johnson, London had lost 17% of permeable surfaces over the last 40 years, and without action, England's capital could be vulnerable to flooding even during normal rainfall.[2] The Greater London Authority (GLA) has allocated over £13 million to green infrastructure projects supporting climate adaptation since 2016, with £1.8 million directed specifically toward surface water flood risk mitigation through SuDS integration.[3] In addition, the government also finances a variety of private entities through public subsidies to provide the desired level of GSI in new development. Although no single funding stream covers the entire cost, these investments from various channels have supported hundreds of discrete projects incorporating rain gardens, swales, green roofs, and permeable paving systems.

Beyond the GLA, several other public entities contribute to SuDS funding, including Transport for London, Thames Water, central government departments, and local borough authorities, alongside various partnership arrangements and contracting-out arrangements. Across all of these channels, conservative estimates indicate that at least £30-40 million in public funding has been allocated since 2012. Appendix Table E.1 provides details on the funding structure of the program, including grant source, allocation method, cost coverage, and government level. To alleviate fiscal pressure on environmental budgets, some SuDS projects at the local level are subsidized or delivered through mandated private provision, with the government providing incentives or mandating private actors to provide an optimal level of investment. For example, the SuDS program has been incorporated into the urban planning frameworks in *The London Plan 2021*, which mandates SuDS integration into new development projects. As a planning requirement, developers must provide SuDS in major developments in critical areas (*Policy 5.13 Sustainable Drainage*) which states that developers should utilize SuDS unless there are practical reasons for not doing so. However, there is still substantial demand for public subsidies to retrofit existing gray infrastructure and fund scheduled routine maintenance in areas without new development, especially in central areas.

Before embarking on my empirical analysis of the SuDS program, it is worth describing the institutional context surrounding how these GSI sites are selected. There are two distinct delivery channels with different selection logics. For planning-led new development, London boroughs, as both Local Planning Authorities and Lead Local Flood Authorities (LLFAs), have been required to secure SuDS on *all* major planning applications since April 2015 (GLA, 2015). This is governed by *The London Plan Policy 5.13*, which (i) requires developers to use SuDS unless impracticable, (ii) aims for greenfield run-off rates, and (iii) sets a drainage hierarchy that prefers surface or visible measures over subsurface storage where possible. Boroughs are also instructed to use their Surface Water Management Plans (SWMPs) to identify problem areas and develop policy responses. In short, for new builds, SuDS siting follows the planning pipeline (where large schemes are happening). Therefore there is generally no extra selection process besides the development scale and site characteristics of the project itself.

For existing neighborhoods, the GLA advances an opportunistic, multi-partner retrofit strategy that "piggybacks" on planned capital and maintenance programs across sectors. The Action Plan instructs the GLA and partners to establish each sector's future capital and revenue programs and identify opportunities for simultaneous sustainable drainage retrofit. This is spelled out for housing repairs, highway resurfacing, hospital works, utilities' capital plans, and more. In practice, retrofit timing and siting depend on where works are already scheduled and where co-funding windows open up. That said, projects are commonly woven into broader objectives of upgrading facilities or improving the appearance and function of buildings and public spaces. Delivery may coincide with regeneration or public-realm works, although borough risk evidence and drainage constraints are the primary targeting criteria of the program. Accordingly, because site selection is institutionally mediated rather than as-if random, I complement the research design by first auditing pre-treatment covariate balance in Section 2.3.

2.2. Data

The Drain London Partnership provides a SuDS retrofit map containing spatial information on all SuDS installations implemented across London. Each geocoded entry includes coordinate location, SuDS typology (rain gardens, permeable paving, green roofs, or bioretention systems), the delivery year, and the implementing organization. The database encompasses 647 individual installations across London's 33 local authorities, though reporting granularity varies substantially by borough. Some authorities report each rain garden as a separate observation, while others aggregate multiple installations within a street or development as a single entry. To address this inconsistency, I standardize observations by aggregating installations within 50-meter buffers and classifying them by primary SuDS type, supplemented by field surveys to resolve ambiguous cases.[4] As shown in Fig. 1, the program was gradually rolled out and proliferated across Greater London since 2012 till now, the newly implemented information is updated annually in the retrofit map.

My analysis combines multiple administrative and environmental datasets to explore the determinants of GSI treatment and factors associated with its impacts. I use Layer Super Output Areas (LSOAs), the smallest unit in UK census statistics, as the unit of analysis and match each installation point to its LSOA treatment area. This allows me to compare the baseline characteristics by treatment status before the program implementation and to control for fixed effects. To identify

[2] See the November 23, 2015 *Stormwater Report* Article, "SuDS Plans Unfold In London".

[3] See *London Sustainable Drainage Action Plan*, by Greater London Authority (October 2015). The latest funding and project information can be accessed on the program website: https://www.london.gov.uk/programmes-and-strategies/environment-and-climate-change/climate-change/surface-water/london-sustainable-drainage-action-plan, (accessed May 27, 2025).

[4] For the 12 percent of installations featuring multiple integrated components, I assign the classification based on the dominant element by surface area. Sensitivity analysis using alternative classification schemes yields qualitatively similar results.

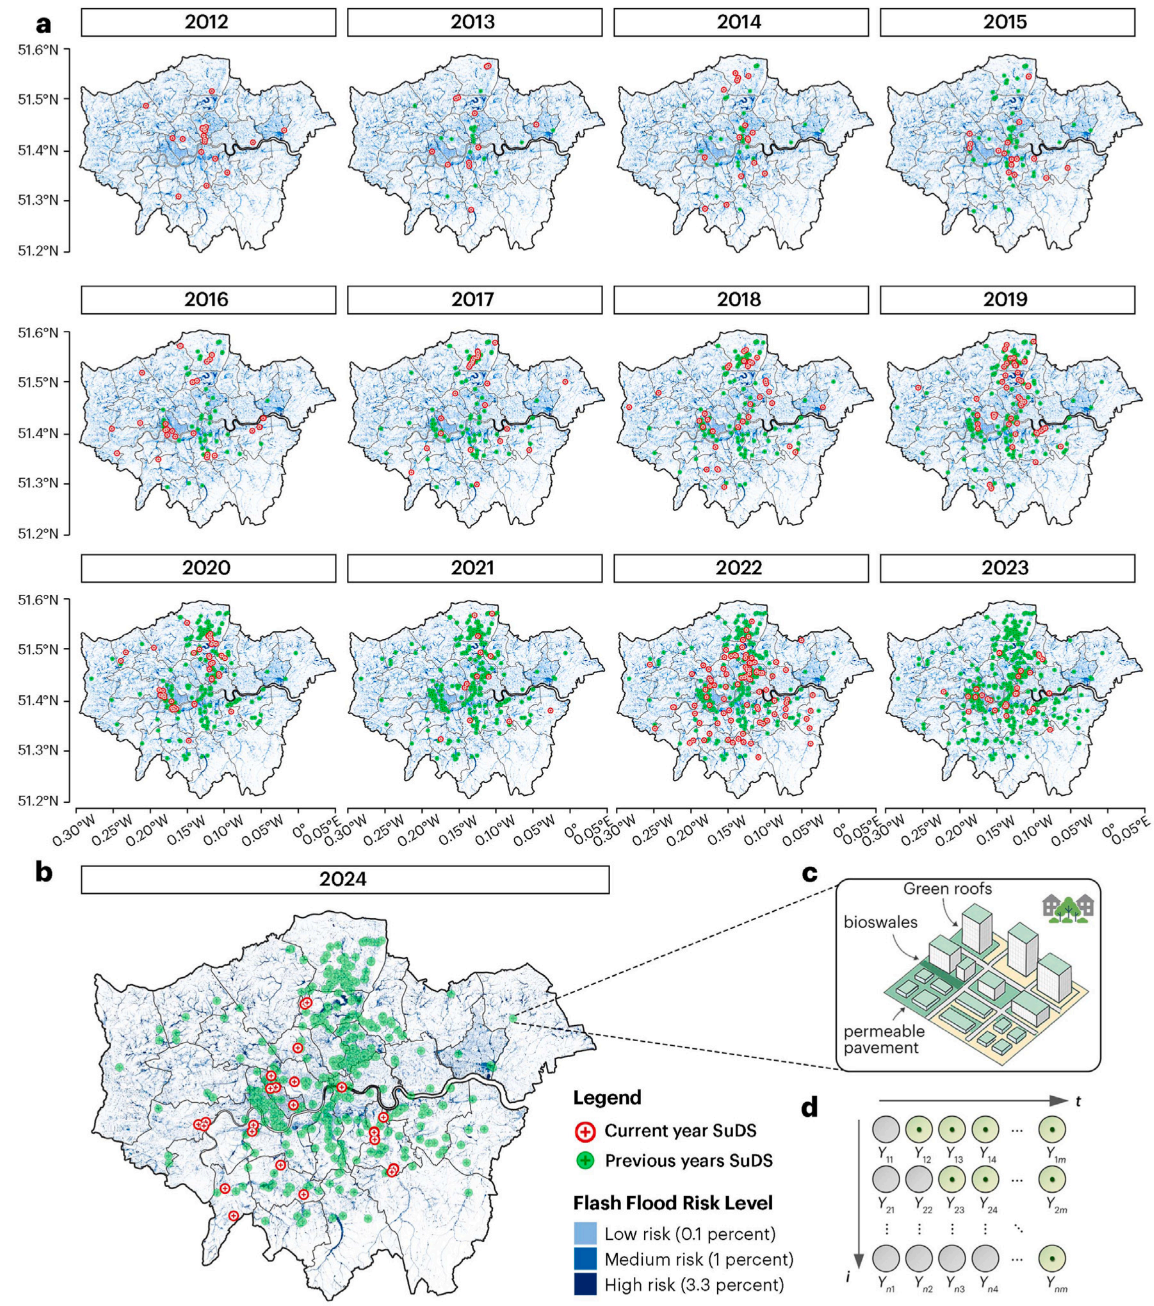

Fig. 1. Rollout of SuDS program in Greater London,2012–2024.
Notes: Panel (a) displays the annual rollout of Sustainable Drainage Systems (SuDS) across Greater London's 33 local authorities from 2012 to 2023. Red and green circles indicate new and previous SuDS deliveries in each year. The base map shows flash flood risk levels from the Environment Agency's Risk of Flooding from Surface Water (RoFSW) model, with darker shades indicating higher annual exceedance probabilities: light blue (0.1–1 percent), medium blue (1–3.3 percent), and dark blue (greater than 3.3 percent). Panel (b) provides the complete spatial distribution as of 2024, with boundaries demarcating 33 local authorities. Panel (c) illustrates the primary SuDS typologies implemented: green roofs, bioswales, and permeable pavement. Panel (d) presents a schematic representation of the staggered adoption pattern, where Y_{it} denotes SuDS presence in location i at time t, with implementation proceeding asynchronously across neighborhoods.

the vulnerability channel and heterogeneous treatment effects, I use the Risk of Flooding from Surface Water (RoFSW) dataset from the UK Environment Agency (EA). The RoFSW employs a probabilistic modeling approach, combining outputs from local topography, rainfall patterns, and historical data to generate flood depth predictions at 2-meter resolution under three annual exceedance probability (AEPs): 3.3 percent (1 in 30 years, high risk), 1 percent (1 in 100 years, medium risk), and 0.1 percent (1 in 1000 years, low risk).

To summarize area-level surface-water exposure in a single intensive statistic, I construct a flood-probability index (FPI) following (Hallegatte et al., 2013) for cross-sectional comparisons of flood vulnerability.[5] I then classify LSOAs into *Low*, *Moderate*, or *High* categories based on the presence of non-trivial surface-water hazard as defined by EA thresholds. I assign an LSOA to the high-risk group if at least

[5] For each LSOA, I compute $FPI_i = (0.033 \times Area\_H_i + 0.010 \times Area\_M_i + 0.001 \times Area\_L_i)/Total\_Area_i$, where $Area\_H_i$, $Area\_M_i$, $Area\_L_i$ denote the areas falling in the high-, medium-, and low-annual exceedance probability and $Total\_Area_i$ is total LSOA area. The index is thus the LSOA's expected share of land inundated in a randomly drawn year.

10 percent of its land area falls within the high-AEP zone (3.3%), to the moderate-risk group if high-AEP exposure is below 10 percent but at least 10 percent of its land lies in the medium-AEP zone (1–3.3%), and to the low-risk group if less than 10 percent of its land is in either the high or medium bands. It bears noting that RoFSW differs from the EA's traditional flood zones — which assess fluvial and coastal flooding — by specifically modeling surface water flooding. This surface water-specific assessment provides superior granularity relative to conventional flood zone mapping, capturing localized flooding patterns that are essential for evaluating spatially targeted drainage interventions. While the EA periodically updates the RoFSW methodology to incorporate improved local modeling, the underlying physical flood risk remains largely stable over my study period, as topographic and hydrological conditions that determine surface water accumulation evolve only over decadal timescales. I use the 2013 baseline risk assessment, which predates our treatment period and is by no means affected by subsequent SuDS implementation.

I also employ Sentinel-1 synthetic aperture radar (SAR) imagery at 10 m resolution obtained from the European Space Agency (ESA) to independently verify surface-water flooding extent during the July 2021 flash-flooding event in London.[6] Sentinel-1 SAR data have the distinct advantage of capturing surface conditions irrespective of cloud interference and are effective for both daytime and nighttime observations. I specifically utilize the VH (vertical transmit, horizontal receive) polarization channel of Sentinel-1 SAR, which is widely recognized as an effective indicator for inundation mapping due to its sensitivity to shallow water surfaces (Clement et al., 2018). To isolate urban road flooding relevant to my analysis and avoid noise from other land cover classes, I apply a spatial mask defined as a 10-meter buffer around the full network of roadways derived from OpenStreetMap. Using this mask, I document temporal changes in VH backscatter values as direct evidence of the spatial and temporal extent of surface-water flooding along urban roadways.

I obtain property-level information from HM Land Registry's Price Paid Data (PPD), which records the universe of residential property transactions in England. The dataset contains 0.97 million transactions over my sample period of 2010–2024 in London, with each record including the transaction price, completion date, full address fields, property type (detached, semi-detached, terraced, or flat), construction vintage (new build versus established), and tenure (freehold or leasehold exceeding seven years). I geocode each transaction using the postcode and Primary Addressable Object Name to assign each property to the corresponding LSOA. To address the temporal lag between transaction completion and registration, which typically ranges from two weeks to two months, I use the transfer deed date rather than the registration date and exclude the most recent two months from my analysis. Besides, I restrict the sample to Category A transactions, representing standard residential sales, thereby excluding repossessions, buy-to-let purchases identifiable through mortgage data, and transfers to non-private entities that entered the dataset in October 2013. I also assess whether SuDS capitalization differs in non-arm's-length transactions by re-estimating the baseline specification on the full PPD (Category A and B) sample and interacting the treatment effect with an indicator for non-standard transfers (see Appendix Table E.2).

2.3. Baseline characteristics by treatment

In Table 1 column (1) to (3), I present summary statistics of the baseline characteristics of the 4994 LSOAs in Greater London by GSI treatment status. All data are collected from the 2011 census, which turns out to be an *ex-ante* measure that predates SuDS program implementation and helps prevent possible "post-treatment bias" (Abadie, 2005). Column (4) reports the differences in means between the treatment and control groups. Relative to the control group, the SuDS treatment group was more likely to include underrepresented minorities (Black, African, Caribbean, and Black British) and to have a larger share of early-career working-age residents (16–29). On the other hand, the treatment group has a similar retired population and a lower share of nuclear households. However, the proportion of White residents, the share of senior working-age groups (30–65), as well as total population and population density, appear to be unrelated to treatment status.

With respect to the socioeconomic status prior to actual treatment, SuDS-adopting neighborhoods tend to have 6.5 percent more renter households and are characterized by significantly lower homeownership rates. Residents in the treatment group are slightly more likely to have lower educational attainment on average, though there is no statistically significant evidence that they earn less compared to the control group. The primary differences between the two groups are in transportation: automobile ownership is 0.14 vehicles per household lower in the treatment group, and the probability of very poor public transit access (PTAL 0–1) is 11 percentage points lower, offset by the same amount of increase in the probability of good public access (PTAL 4–6). This implies a SuDS policy preference for enhancing public-transport connectivity while placing less emphasis on private mobility.

Taken together, neighborhoods selected for the SuDS program tend to be younger, more ethnically diverse, more public transport oriented, and socioeconomically disadvantaged, although most of the baseline characteristics are relatively small in magnitude. Differences in observable characteristics yield identification concerns for cross-sectional comparisons, which would confound program effects with pre-existing disadvantage. Fig. B.1 tracks and compares trends in annual housing transaction prices for neighborhoods with and without GSI treatment over time (see Appendix B for a more detailed description). Nevertheless, the raw differences between these two groups do not isolate the treatment effect because baseline socioeconomic characters differ. Moreover, GSI might be introduced in areas with poor gray infrastructure in response to higher flooding risk, which could directly affect the real estate market and other economic activities. These selection biases motivate the use of a difference-in-differences (DiD) design that accommodates baseline differences, as long as outcomes would have trended in parallel in the absence of policy change (Black et al., 2023).

To further defend against these confounding factors, I also generate comparison groups, via coarsened exact matching (CEM) and inverse propensity score weighting, which are similar to the treated group with respect to these prior covariates. More precisely, the matched and weighted groups are balanced across all covariates as shown in Fig. C.2, such that any remaining difference between treated and control groups can plausibly be attributed to the effect of GSI treatment rather than preexisting differences. Details of this procedure are described in Appendix C.

3. Empirical framework

To compare the economic outcomes between LSOAs that have been assigned GSI as an adaptation measure to mitigate chronic flooding risk and those that have not, I begin by augmenting a standard hedonic pricing model to include a dummy variable indicating whether the site has at least one GSI treatment. Let D_s represent a dummy variable that equals one if the sth LSOA has GSI and zero otherwise. Following the hedonic price model (Rosen, 1974), I first consider the outcome to be the transaction price of property i in sth LSOA and use the following log-linear specification:

$$\ln P_{i,s} = \alpha + \beta D_s + \boldsymbol{X}_i^\top \boldsymbol{\gamma} + \boldsymbol{Z}_s^\top \boldsymbol{\varphi} + \epsilon_{i,s}, \tag{1}$$

[6] Most Sentinel-1 SAR images between July 4 and July 31 in 2021 cover the entire Greater London region; in cases where spatial coverage is incomplete, each scene nonetheless exceeds 70 percent coverage, and I restrict my analysis to neighborhoods with complete data.

Table 1
Baseline neighborhood characteristics by SuDS program assignment.

Baseline characteristics	Means			Two-sample t-test	
	All LSOA (1)	LSOA with SuDS (2)	LSOA without SuDS (3)	Mean difference (4)	Standard error (5)
A. Demographic					
Race = White	0.605	0.620	0.604	0.016	(0.009)
Race = URM	0.131	0.161	0.128	0.032***	(0.006)
Age = 16–29	0.221	0.230	0.220	0.009**	(0.003)
Age = 30–64	0.468	0.467	0.472	0.004	(0.002)
Age = 65+	0.113	0.115	0.113	0.002	(0.004)
Nuclear household	0.418	0.380	0.421	−0.040 ***	(0.006)
Total population	1697	1698	1691	−6.62	(13.6)
Population density	97.4	100.0	97.2	3.01	(3.21)
B. Socio-economics					
Homeownership	0.220	0.185	0.222	−0.037 ***	(0.006)
Social rented	0.235	0.295	0.230	0.065***	(0.011)
Unemployment	0.074	0.081	0.073	0.007***	(0.002)
Qualification ≤ High school	0.373	0.396	0.371	0.024**	(0.008)
Qualification = University	0.098	0.096	0.097	−0.001	(0.001)
Median household income (£)	35,756	36,017	35,735	282	(687)
C. Transportation					
Car per household	0.848	0.715	0.859	−0.144 ***	(0.018)
PTAL = 0–1 (poor access)	0.238	0.134	0.247	−0.113 ***	(0.013)
PTAL = 4–6 (good access)	0.282	0.386	0.273	0.113***	(0.022)
Wilks' Lambda joint F-test	–	–	–	9.451***	(0.967)
Observations (N)	4835	368	4467	–	–

Notes: The table summarizes baseline year (2011) characteristics of Lower Super Output Area (LSOA) in Greater London, separately for those with SuDS treatment at the end of the research period (2024) and never receive treatment till then. These statistics are obtained from 2011 Census Statistics for England and Wales. Column (4) reports the difference in mean of the treatment group and control group, with the standard error shown in column (5). PTAL = Public Transport Accessibility Levels assessed by Transport for London (TfL). Each area is graded between 0 and 6, ranging from a very poor access to public transport to excellent access. URM = Underrepresented Ethnic Minority, specifically Black, African, Caribbean, and Black British. The penultimate row presents Wilks' Lambda F-statistic for joint significance of all differences by treatment. Standard error are in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

where $\boldsymbol{X}_i^\top$ is a vector of property-level controls including year built, housing type (i.e. detached, semi-detached, terrace, or flat),[7] and ownership (leasehold or freehold). $\boldsymbol{Z}_s^\top$ is a set of location-based neighborhood characteristics that vary across LSOAs. In principle, $\boldsymbol{Z}_s^\top$ is largely unobservable that represents general environmental and location components; these omitted features therefore appear in the error term.

Since I am not interested in predicting housing prices or local vitality at any specific point in time, estimating the coefficients for the unobserved variable $\boldsymbol{Z}_s^\top$ has little value and is extremely difficult. Instead, I focus on the coefficient on the causal variable of interest D_s, which captures the change in housing prices relative to the untreated counterfactual. A canonical two-period difference-in-differences (DiD) estimator is helpful for identifying this effect (Card and Krueger, 2000), so that I can estimate average treatment effects (ATE) from cross-sectional data as follows:

$$\beta_{DiD}^{2\times2} = [E(Y_s^{Post}|S_i = 1) - E(Y_s^{Pre}|S_i = 1)] - [E(Y_s^{Post}|S_i = 0) - E(Y_s^{Pre}|S_i = 0)] \tag{2}$$

The key idea of the simple 2×2 DiD estimator is to difference out the unobserved time-invariant effects before and after the intervention. Their invisibility notwithstanding, assuming they are fixed allows us to control for them. The method works well in cross-sectional data; however, an annual panel, such as housing transaction records, complicates the analysis and diverges from the 2×2 setup. Moreover, London's SuDS program rolls out gradually from 2015 to the end of our study period, which means treatment timing varies. Under these circumstances, we cannot rely on a canonical DiD estimator because the pre-post dummies are not clearly defined for all treatment observations (De Chaisemartin and d'Haultfoeuille, 2023). Therefore, I aggregate housing transaction data on an annual basis—this multi-year panel enables a two-way fixed effects (TWFE) model (Carpenter and Dobkin, 2011) to compare longitudinal outcomes across LSOAs experiencing different evolutions of exposure to treatment. Specifically, I estimate SuDS treatment effect using a TWFE model that includes both LSOA fixed effects and year fixed effects:

$$\ln Y_{i,s,t} = \alpha_s + \lambda_t + \tau D_{s,t} + \boldsymbol{X}_i^\top \boldsymbol{\gamma} + \epsilon_{i,s,t} \tag{3}$$

where α_s is the unit fixed effects and captures unobserved, time-invariant characteristics of each LSOA, λ_t is the year fixed effects and captures common, atypical year-to-year variation; and $D_{s,t}$ is a dummy variable indicating whether the sth LSOA has received GSI by year t (equal to one starting in the first treatment year and in all subsequent years). The treatment effect of interest is τ.

The regressions are weighted by LSOA population from the 2021 census, with standard errors clustered at the LSOA level. According to the DiD decomposition theorem proposed by Goodman-Bacon (2021), the treatment timing structure of SuDS allows us to conceptually distinguish an early-treatment group K ($S_{i,t}$ turns to 1 at $t = k$), a late treatment group L ($S_{i,t}$ turns to 1 at $t = l$), and a consistently untreated group U. As depicted in Fig. 2, I set k as an earlier SuDS delivery year (e.g., 2012) and l to a later delivery year (e.g., 2021). Assuming that the effect of SuDS is a one-time, discrete improvement, Panel A and B show 2×2 switcher-control comparisons in which the TWFE estimator $\hat{\tau}_{KU}^{2\times2}$ and $\hat{\tau}_{LU}^{2\times2}$ simply reduce to canonical DiD cases.

The validity of the TWFE estimand rests on two assumptions (Sun and Abraham, 2021): (1) *Parallel trends*, which posits that $E[\ln Y_{i,t}(0)|U] - \ln Y_{i,t}(0)|\{K, L\}]$ evolves identically in the pre- and post-windows. In other words, in the absence of treatment, the average

[7] London's housing stock is dominated by flats, particularly in inner boroughs, and these units are generally less exposed to surface-water flooding than terraced or detached houses. Therefore, housing composition could confound estimated capitalization effects if treated and untreated areas differ systematically in property types.

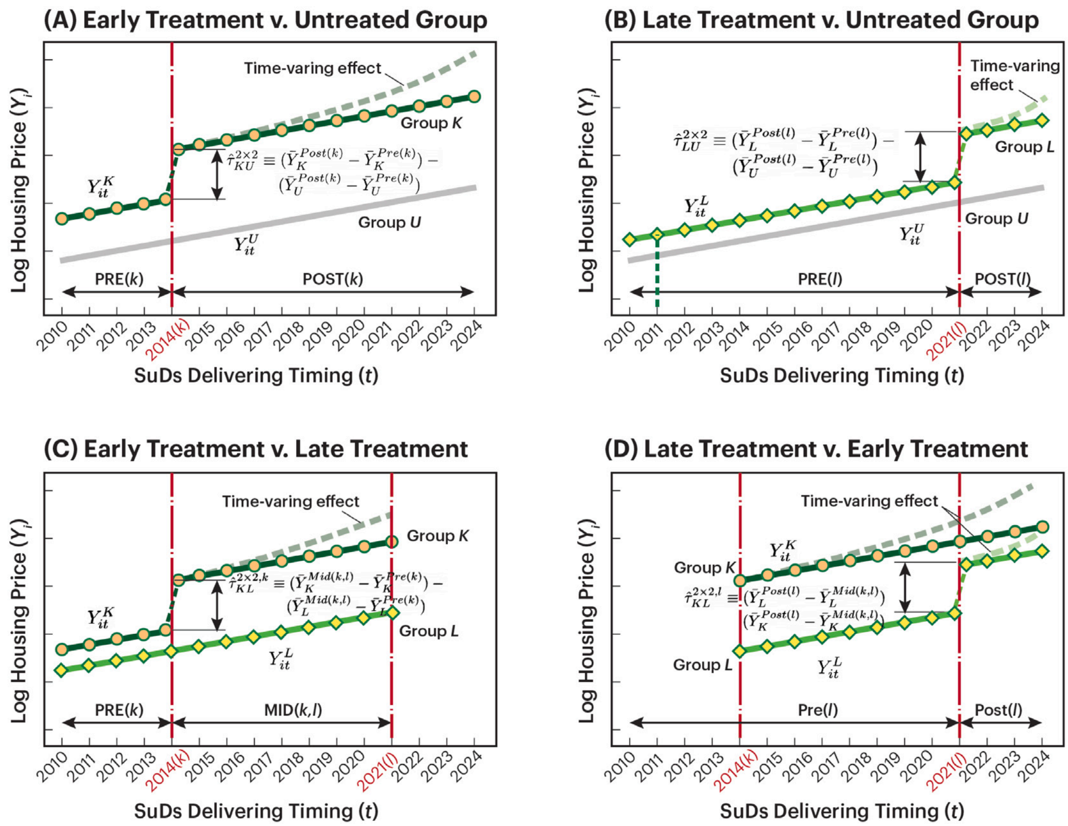

Fig. 2. Decomposition of SuDS average treatment effect from two-way fixed effects model.
Notes: Panels A through D illustrate the four two-by-two comparisons that form the Bacon decomposition (Goodman-Bacon, 2021) of a two-way fixed effects estimator when SuDS program rolls out since 2014. Three groups are compared: an "early" treatment group initiating in 2014 (Group K), a "late" treatment group initiating in 2021 (Group L), and an untreated group (Group U). The vertical dashed lines mark first-treatment dates for each group. In each panel, the slope of the green dashed line indicates a potential time-varying or cumulative treatment effect. The standard 2×2 difference-in-differences (DiD) estimator is valid under a one-off treatment-effects assumption, as shown as the gaps under each scenario. The estimations in the second row would be biased if there is a cumulative treatment effect that can complicate inferences from a single-coefficient two-way fixed effects regression.

outcome of the treated group would have changed in the same way as the average outcome of the non-treated group. (2) *Homogeneous or constant treatment effect*, which restrict the treatment's impact to the jump-off window. If the treatment effects grow over time, the early group K cannot be used as a legitimate "control" for the late group L. Without these two assumptions, the estimand τ need not describe a causal effect.

It is feasible to assess the parallel trends by leveraging event study specifications and conducting a joint-F test on the lead coefficients (Deryugina and Molitor, 2020), but the assumption is untestable because it is a statement about a counterfactual that cannot be definitively proven.[8] Fortunately, my data include a rich set of pre-treatment covariates for each LSOA, which makes it possible to condition on various differences between treatment and control groups, thereby reducing residual selection on observables when overlap between treated and not-yet-treated LSOAs is limited. I thereby complement the canonical DiD with two robustness designs, including a matching-then-DiD estimator that pre-processes the panel and an inverse-probability-weighted DiD (IPW-DiD) that reweights controls to reproduce the treated covariate distribution. It is possible to use both checks simultaneously in an effort to rid the estimates of bias; this approach is described in detail in Appendix C.

The constant-effect assumption is likely to be violated in my setting, as green infrastructure tends to exhibit time lags in generating environmental and economic impacts, though the extent and direction of this heterogeneity require a judgment call (Bae et al., 2021). Under these circumstances, τ in Eq. (3) not only captures DiD comparisons between switchers and controls; the Frisch–Waugh–Lovell theorem tells that all available variation in S_{it} is used to calculate the regression coefficient.[9] Because post-treatment outcomes of the early cohorts can differ for many reasons besides serving as a control, some of the TWFE comparison weights turn negative (Panel D, Fig. 2). The resulting "contaminated" estimate of τ averages heterogeneous effects with the wrong signs and the wrong timing. To deal with negative-weighting bias from unintended comparisons, I adopt a modern event study specification that follows the imputation method of Borusyak et al. (2024). It nests the original TWFE estimator as the special case $k = 0$ and makes every step, including counterfactual imputation, explicit.

4. Results

4.1. Flash flooding discount effect

I first present results for the vulnerability channel, examining whether chronic flash flooding has negative impacts on housing price in London, consistent with existing literature. I augment a standard

[8] Although the common trend assumption is inherently untestable, the results from a joint F-test of the leads still generate testable implications: specifically, outcome trends should be similar before the policy change. A detailed discussion on this test and results is presented in Section 4.4.

[9] Under Bacon's decomposition, the estimand τ is a weighted average of all possible 2×2 DiD estimators, as shown in Fig. 2.

Table 2
HPM estimates of the discount effects of flash–flooding risk on housing price.

	2010 Sample				2024 Sample			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
High flash–flooding probability	−0.202***	−0.239***	−0.150***	−0.183***	−0.151***	−0.175***	−0.121***	−0.132***
(1 in 30 years, 3.3%)	(0.032)	(0.032)	(0.028)	(0.029)	(0.039)	(0.037)	(0.035)	(0.035)
Moderate flash–flooding probability				−0.040***				−0.102***
(1 in 100 years, 1%)				(0.006)				(0.006)
Low flash–flooding probability				−0.024***				−0.042***
(1 in 1000 years, 0.1%)				(0.008)				(0.011)
Type = Detached		0.252***	0.261***	0.261***		0.586***	0.822***	0.820***
		(0.052)	(0.046)	(0.046)		(0.054)	(0.050)	(0.050)
Type = Semi-Detached		0.191***	0.203***	0.203***		0.053	0.337***	0.336***
		(0.052)	(0.046)	(0.046)		(0.048)	(0.045)	(0.045)
Type = Terraced		0.174***	0.181***	0.181***		−0.085**	0.137***	0.137***
		(0.052)	(0.046)	(0.046)		(0.042)	(0.044)	(0.044)
Type = Flats/Maisonettes		0.149***	0.144***	0.144***		−0.133***	−0.025	−0.025
		(0.052)	(0.046)	(0.046)		(0.044)	(0.042)	(0.042)
Newly built property		0.009	−0.113***	−0.113***		−0.347***	−0.458***	−0.486***
		(0.006)	(0.006)	(0.006)		(0.013)	(0.011)	(0.012)
Leasehold		−0.087***	−0.209***	−0.209***		−0.461***	−0.514***	−0.503***
		(0.020)	(0.0158)	(0.015)		(0.033)	(0.031)	(0.030)
Standard price-paid entry		0.293	0.186	0.015		0.274***	0.255***	0.242***
		(0.204)	(0.153)	(0.154)		(0.006)	(0.006)	(0.006)
Borough fixed effects	No	No	Yes	Yes	No	Yes	Yes	Yes
Property-level controls	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Granular risk band dummies	No	No	No	Yes	No	No	No	Yes
Observations (N)	58,243	58,243	58,243	58,243	49,035	49,035	49,035	49,035
Adjusted R^2	0.000	0.126	0.499	0.500	0.000	0.149	0.344	0.345

Notes: This table reports HPM estimates of the discount effects associated with flash flooding risk on residential property values in 2010 and 2024 samples. The dependent variable is the natural logarithm of transaction price. Flash flooding risk indicators are constructed from the Environment Agency's RoFSW probabilistic model, with properties classified into high (>3.3 percent annual probability), medium (1–3.3 percent), or low (0.1–1 percent) risk categories based on the maximum predicted flood depth within a 30 m buffer. The omitted category comprises properties in non-risk areas. Property-level controls include categorical variables for property type, construction vintage, tenure and transaction type. Borough fixed effects control for time-invariant spatial heterogeneity across London's 33 local authority districts. Granular risk band specifications in columns (4) and (8) include all three risk level indicators simultaneously. Sample sizes reflect properties with complete covariate information excluding the most recent two months to account for registration lags. Robust standard errors clustered at the LSOA level appear in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

hedonic price model in log-linear specification and include a series of dummy variables indicating the average probability of flash flooding surrounding the property[10] from the RoFSW probabilistic model. While RoFSW probabilities offer a consistent and spatially detailed measure of underlying flash-flood exposure, this hedonic exercise does not incorporate the timing or severity of realized flood events. Table 3 reports results on the impacts of flash flooding on housing prices in two cross-sectional samples, with varying sets of covariates, following the path analysis in causal inference (Angrist and Pischke, 2009). Column (1) and (4) provide the raw difference in log prices between residential properties that fall into high-flash-flood risk areas (i.e. one major flooding event predicated every 30 years) and those outside the risk area. Not surprisingly, these raw gaps imply substantial discounts of 20.2% and 15.1% in 2010 and 2024 respectively.

One may suspect that some of these gaps are almost certainly due to selection bias, because flash flood chance is not only determined by natural random factors like climatology and typology, but also the gray infrastructure and impervious surface, it is therefore of interest to add property-level controls and spatial fixed effects in my model. In column (2) and (5), I include property-level control dummies. Controlling for property characteristics yields even larger differentials (23.9% and 17.5%) than the raw comparisons. When borough fixed effects are added, the estimated association between the risk indicator and housing prices falls back to 15.0% and 12.1%, with a substantial increase in goodness of fit (adjusted R^2), as reported in columns (3) and (6) of the table. Although the magnitude of the estimate fluctuates with the inclusion of these covariates, the flood-risk indicator continues to show a sizeable and statistically significant discount in housing transaction prices, with t-ratios of 5.26 in 2010 and 3.42 in the 2024 sample.

RoFSW also provides medium- and low-risk levels in its probabilistic model. To explore differential discount effects associated with different risk bands, I break down the risk dummy into three risk-level indicators. These results are reported in column (4) and (8): for properties with a 1% flash-flood risk, the discount shrinks to only 4.0 percent, remarkably lower than in high-risk areas in 2010. In contrast, the discount rates for high- and medium-risk areas become similar in 2024. It could be interpreted as the fact that households are more sensitive to medium-level flooding risk amid growing concern about climate uncertainty. The price effect associated with low-probability flash-flood risk remains small in magnitude in both years (2.4% and 4.2%), though it is statistically significantly different from zero. These estimates are consistent with the meta-analysis of flood zone price discount range in Contat et al. (2024), but lower in magnitude than Skouralis et al. (2024), who use only single-family housing transactions (which are more sensitive to flood risk) across England and not restricted to urban flash flooding.

The results in Table 2 show the discount associated with high flooding risk in 2024 is 3–5 percent lower than the year 2010 when there was no systematic adaptation measure. It is reasonable to speculate that this attenuation in depreciation reflects an adaptation effect. However, since any other climate adaptation measures and risk-information communications could also affect the housing-market outcomes, these results should be interpreted with caution. The remainder of my analysis focuses directly on GSI investments as an external intervention to examine this speculation.

[10] I use a 30-meter radius to construct buffer area around each property and summarize the flooding chance information in this buffer area. Estimated effects are similar for different buffer radius like 40 or 50 m.

4.2. Identifying average SuDS treatment effect on housing price

To quantify the magnitude of these treatment effects and move beyond suggestive graphical evidence, Table 3 presents my main difference-in-differences estimates of the SuDS program's impact on housing markets. The empirical specification follows the two-way fixed effects model in Eq. (3), exploiting the rollout of green infrastructure investments across London's neighborhoods to identify causal effects on both property values and market liquidity. Throughout, the identifying variation comes from the staggered timing of first SuDS delivery within a 2010–2024 panel of 981,250 property transaction entries in 4835 LSOAs. This causal attribution rests on a common-trends assumption, which I test in Section 4.4, and does not rely on covariates balance. However, as discussed in Appendix C, such a pre-trends claim could be strengthened when covariates are balanced or accounted for. To maximize the internal validity of the study, I present additional results from two robustness designs — matching-then-DiD and inverse-probability-weighted DiD — in Appendix Table C.1, which are in general very similar to my baseline DiD estimates.

Panel A examines the program's capitalization into housing prices. The simplest specification in column (1) indicates that SuDS implementation generates a statistically significant 3.7 percent increase in property values. The magnitude is consistent with broader climate adaptation appreciation effects in the recent meta-analysis (Wu et al., 2026) and the distance decay estimates for Philadelphia GSI in Cohen et al. (2025), which show effects attenuating from 10 percent at 25 m to 3.2 percent at 50 m.[11] This effect persists across alternative specifications that progressively address potential confounders. The inclusion of year-quarter fixed effects in column (2), to absorb seasonal variation in London's housing market, leaves the estimate virtually unchanged at 3.8 percent. In my preferred specification, when I weight observations by LSOA population in column (3) to account for the differential density of transactions across neighborhoods, the effect moderates slightly to 3.2 percent but remains highly significant. Most stringently, column (4) augments the baseline model with LSOA-specific linear time trends to control for any differential pre-existing trajectories across treated and control neighborhoods—a particularly demanding test given my relatively short panel. Even under this conservative specification, I still find a robust 3.5 percent price appreciation attributable to green infrastructure investment. These point estimates provide evidence that SuDS appreciation effect is neither an artifact of intra-annual housing cycles nor driven by differential pre-treatment trajectories.

Next, I consider whether exposure to flash-flood risk affects the GSI adaptation treatment effect. Even though funding for subsidizing SuDS programs was not directly targeted at frequently flooded areas, I would expect the treatment effect to be more salient in those hazardous areas. In doing so, I stratify the sample into high-, moderate-, and low-risk categories (Section 2.2) and re-estimate average treatment effects in these subsamples separately. Perhaps unsurprisingly, in high-risk areas with an annual flooding probability greater than 3.3 percent, the treatment effect ranges from 3–5 percent across specifications, with the population-weighted estimate suggesting a pronounced 5.2 percent appreciation.

The results align with the theoretical prediction that climate adaptation investments should generate larger benefits where baseline vulnerability is highest. The coefficients are not sensitive to the inclusion of quarter fixed effects and population weights, but they attenuate and turn to be marginally significant when more restrictive LSOA-specific trends are included. The modest pre-existing differences in price trajectories are partly due to rising insurance premiums and heightened awareness of flood hazard in high-risk neighborhoods. Once a tract-specific linear trend absorbs the secular decline, the remaining level impact of the first installation is mechanically smaller. I also estimate SuDS effects separately for more flood-exposed dwelling types (terraced and detached houses) and for flats, which are typically less directly exposed. The estimated capitalization effects, as shown in Appendix Table E.4, are substantially larger for houses, consistent with greater benefits where baseline flood risk is higher. This gap is nonetheless consistent with a smaller but still meaningful role for flood risk in flat valuations, since leaseholders are typically not insulated from flood-related costs even when their individual unit is not directly inundated.[12]

Properties in moderate risk LSOAs, on the contrary, exhibit no statistically significant price response to SuDS implementation. The point estimates show non-trivial positive treatment effects between 1.1 and 1.5 percent, but with standard errors that preclude rejection of the null hypothesis. This is particularly informative when considered alongside my earlier hedonic estimates showing that medium flood risk still commands a 4–10 percent price discount in London's housing market. The absence of significant capitalization suggests that in areas with moderate flooding probability (1–3.3 percent annually), the perceived benefits of distributed green infrastructure may be insufficient to measurably shift property valuations. This could reflect either genuinely limited flood reduction efficacy at intermediate risk levels or households' imperfect information about GSI adaptation measures in the absence of salient flooding events.

Finally, there is a salient price appreciation in low-risk areas, with effects ranging from 3.9 to 4.6 percent and statistical significance across all specifications. It seems counterintuitive that green infrastructure generates larger price premiums in areas with minimal flooding exposure than in the moderate-risk stratum. One possible explanation for this non-monotonic marginal willingness to pay is household risk perception that overweights low-probability catastrophes but underweights intermediate probabilities once insurance, self-protective behavior, or ambiguity aversion are factored in (Kunreuther et al., 2023). However, as shown in Table 2 column (4) and (8), the discounting effects are monotonic by risk stratum in both 2010 and 2024 samples. That ordering is exactly what one would predict from a standard expected-damage model in which buyers value houses according to actuarially fair insurance premia. It seems unlikely that households overreact to small flooding probabilities and underreact to intermediate ones.

Recent research documents non-linear relationships between flood risk, water-related amenities, and housing prices,[13] it is therefore not unexpected to see SuDS capitalization depart from a simple linear mapping from baseline risk to prices. Since the pre-treatment discount schedule is monotone in risk in our setting, the non-monotone pattern of SuDS capitalization reported in Table 3 must originate elsewhere. I

[11] The average property within a treated LSOA lies well beyond 25 m from the nearest installation, hence my neighborhood-level estimates capture a spatially averaged effect consistent with this gradient.

[12] Increases in building-level insurance premia following flood events are commonly passed through to leaseholders via service charges, so flood risk and risk-reducing investments can affect the cost of ownership independent of direct exposure.

[13] For example, Skouralis et al. (2024) show that, once one moves from coarse floodplain indicators to a property-level flood score, medium-risk dwellings can in some cases transact at premia over "zero-risk" properties, plausibly because modest increases in hazard are more than offset by proximity to attractive water amenities. Similarly, Rajapaksa et al. (2017) and Beltrán et al. (2019) report hump-shaped price gradients with respect to river proximity: houses very close to rivers often enjoy net price premia despite elevated exposure, consistent with amenity values outweighing expected damage for a segment of buyers. These findings suggest that the capitalization of water-related risks and benefits is not well approximated by a linear mapping from objective flood probability to price; instead, amenity and risk components interact in ways that can generate non-monotonic patterns across the risk distribution.

Table 3
Effects of green stormwater infrastructure adaptation on housing prices and market activity.

	(1)	(2)	(3)	(4)
Panel A. Dependent variable: log(*Housing Price*), *OLS*				
All LSOA	0.037***	0.038***	0.032***	0.035***
	(0.008)	(0.008)	(0.011)	(0.011)
Observations (N)	981,250	981,250	981,250	981,250
High flooding risk LSOA	0.039***	0.043***	0.052***	0.032*
	(0.016)	(0.016)	(0.019)	(0.018)
Observations (N)	305,003	305,003	305,003	305,003
Moderate flooding risk LSOA	0.017*	0.016*	0.021*	0.023
	(0.009)	(0.009)	(0.012)	(0.015)
Observations (N)	496,515	496,515	496,515	496,515
Low flooding risk LSOA	0.045***	0.046***	0.043**	0.039**
	(0.017)	(0.012)	(0.019)	(0.018)
Observations (N)	179,732	179,732	179,732	179,732
Panel B. Dependent variable: Annual transactions, Poisson pseudo-likelihood estimator				
All LSOA	0.004	−0.002	−0.057	−0.006
	(0.038)	(0.036)	(0.064)	(0.036)
Observations (N)	44,507	165,167	165,167	165,167
High flooding risk LSOA	−0.075	−0.070	−0.160	0.106
	(0.046)	(0.045)	(0.106)	(0.046)
Observations (N)	13,838	51,514	51,514	51,514
Moderate flooding risk LSOA	0.053	0.029	0.034	0.020
	(0.070)	(0.064)	(0.083)	(0.063)
Observations (N)	22,176	82,176	82,176	82,176
Low flooding risk LSOA	0.095	0.103	−0.019	0.107
	(0.084)	(0.083)	(0.162)	(0.083)
Observations (N)	8493	31,477	31,477	31,477
LSOA–Year fixed effects	Yes	Yes	Yes	Yes
Year–Quarter fixed effects	No	Yes	Yes	Yes
LSOA linear time trends	No	No	No	Yes
Population weights	No	No	Yes	No

Notes: Each entry is the coefficient from a separate two-way fixed-effects regression of SuDS treatment on housing-market outcomes. Panel A estimates log transaction price with ordinary least squares at the property–year level and includes property covariates (built year, housing type and ownership). Panel B estimates the annual number of housing transactions in each Lower-layer Super Output Area (LSOA) with the Poisson pseudo–maximum-likelihood estimator. coefficients are shown as percentage changes, transformed in the following way: $100 \times [\exp(\beta) - 1]$. The treatment indicator equals one from the first calendar year in which any SuDS installation is completed within an LSOA and remains one thereafter. LSOAs are stratified into High, Medium, and Low flooding-risk groups based on the Environment Agency's three annual exceedance probability (AEP) bands. All specifications contain LSOA and calendar-year fixed effects; column (2) further absorbs year-by-quarter fixed effects, column (3) applies 2011 Census LSOA population weights, and column (4) augments the baseline model with LSOA-specific linear time trends. Standard errors are estimated using the delta method and are clustered at LSOA level. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

therefore set aside probability-weighting story and turn to the mechanism of green stormwater infrastructure per se. There are various GSI typologies that mix together in London's SuDS practices, like rain garden, bioswale, surface detention basin, green roof, etc. Heterogeneity in treatment effects across risk strata may reflect differences in the engineering effectiveness of GSI typologies in managing flash-flood runoff. On the other hand, different GSI typologies may cover different signals to prospective buyers through their salience (i.e., how easily they convey adaptation information). Homebuyers often use readily observable features as heuristics when valuing homes; low salience can depress willingness to pay even if hydraulic performance is real. I return to these mechanisms later in Section 5.

An additional institutional feature of the UK context is the national flood reinsurance scheme Flood Re, introduced in 2016 to stabilize the private insurance market for high-risk homes by subsidizing flood premiums on properties built before 2009. To explore whether this policy environment shapes the capitalization of SuDS, I augment the baseline two-way fixed-effects model with an interaction between SuDS treatment and a post-2016 indicator. In principle, because subsidized insurance changes the marginal benefit of physical adaptation by smoothing the financial consequences of flood losses, I would expect SuDS capitalization to be somewhat *crowded out* thereafter. The results, reported in Appendix Table E.5, however, show a small but positive shift in the SuDS effect after 2016. A plausible explanation is that any dampening effect of subsidized insurance on the value of physical adaptation is more than offset by the heightened salience of flood risk in the Flood Re era, combined with the fact that many SuDS installations are in post-2009 neighborhoods outside the scheme's eligibility criteria. That said, this test should be viewed as suggestive rather than definitive, as the design relies on a coarse pre-post timing break, and other contemporaneous policy and market changes may also contribute to the observed pattern.[14]

A final interpretation question relates to the extensive-margin treatment indicator I adopt, which switches on when an LSOA first receives any SuDS installation; the retrofit map also allows me to exploit limited variation in the intensity of provision over time. I present two pieces of evidence using an intensive-margin model in which the effect of SuDS is allowed to scale with treatment intensity, supporting the conclusion that my results are not driven by the treatment definition. Specifically, I substitute a continuous treatment variable—the cumulative count of installations per unit area—for the binary indicator in Eq. (3). In a second specification, I adopt a simple categorical dose–response structure that distinguishes between LSOA–years with exactly one SuDS

[14] It is also worth noting that Flood Re is explicitly a transitional scheme, legislated to expire in 2039 with the stated aim of returning the market to risk-reflective pricing thereafter. Insofar as forward-looking buyers price in this sunset provision, Flood Re's dampening effect on the value of physical adaptation should be weaker than that of a permanent insurance arrangement, which may help explain why its interaction with SuDS capitalization is modest in magnitude.

installation and those with two or more, with no SuDS as the omitted category. These results, displayed in Appendix Table E.6 from the second to the last column, are broadly consistent with the baseline estimates, though there is a sizeable drop in the coefficients of the second column when treatment variable is continuous as they stand for per standard-deviation increase in SuDS density. Therefore, the cost–benefit calculations based on the binary treatment remain credible and, if anything, conservative, because they effectively evaluate the program at the average observed intensity.

4.3. Transaction volume effects and market equilibrium

Table 3 Panel B investigates whether SuDS implementation affects housing market liquidity, measured as annual transaction volumes aggregated at the LSOA level. Each coefficient is estimated using Poisson pseudo-maximum likelihood and comes from a separate regression that includes at least LSOA and year two-way fixed effects. The price appreciation documented in Panel A indicates an outward shift in housing demand, reflecting increased willingness to pay for properties in SuDS-treated areas. Under standard assumptions of upward-sloping supply, such demand increases should generate both higher prices and expanded transaction volumes. However, the empirical evidence decisively rejects this prediction. Across all specifications and risk categories, I find no evidence that green infrastructure generates a discernible rise in housing market activity. Point estimates are economically small and statistically indistinguishable from zero, with confidence intervals that rule out substantial effects in either direction.

This null result on the extensive margin is noteworthy for two reasons. First, it suggests that the price appreciation documented in Panel A reflects genuine increases in property values rather than compositional changes driven by selective market participation. If SuDS primarily attracted higher-quality properties to the market or deterred transactions in lower-quality housing stock, we would expect to observe systematic changes in transaction volumes. Second, the absence of liquidity effects alleviates concerns about market frictions or information asymmetries that might prevent efficient capitalization of adaptation benefits. The absence of detectable effects on annual transaction counts is consistent with a setting in which demand shifts outward but housing supply is highly inelastic over the horizon we study. Taken at face value, the combination of higher prices and roughly unchanged quantities is exactly what a simple partial-equilibrium model would predict when supply cannot expand in response to an amenity shock. As I am unable to leverage supply-side data (e.g., new construction permits, planning applications, or new-build completions) to provide direct evidence of a relative decline in new housing supply, my conclusion about supply constraints remains a conjecture, and I leave this endeavor to future work.

Although it is hard to draw strong conclusions from this evidence, several institutional features of London's planning and drainage regime could potentially provide a plausible rationale for such inelasticity. The London Plan 2021, which mandates SuDS integration for new developments in critical drainage areas, imposes additional costs and complexity on residential construction (Chapman and Hall, 2022). Developers must now incorporate permeable surfaces, detention basins, or green roofs into project designs—requirements that increase both upfront capital costs and ongoing maintenance obligations. These regulatory constraints effectively steepen the supply curve in treated areas, as marginal development costs rise with binding sustainability requirements.[15]

Indeed, supply constraints around SuDS are not merely a product of this recent comprehensive planning. Beginning in the mid-2010s, the LSDAP set out a city-wide strategy to embed SuDS across sectors and explicitly framed sustainable drainage as a requirement to be integrated into existing capital and maintenance programs. This strategic push was operationalized through instruments such as the London SuDS Proforma. Additionally, some borough-level planning policies moved earlier towards tighter regulation of surface-water management. For example, Camden began requiring SuDS on all major planning applications in 2014, guided by its Local Plan and Strategic Flood Risk Assessment. Richmond treated sustainable drainage as a material consideration from 2015 onward, insisting that proposals adhere to the London Plan's drainage hierarchy, while Merton required detailed Sustainable Drainage Strategies for major schemes and published its own design and evaluation guidance.

4.4. Event study estimates

Despite robust evidence from my baseline difference-in-differences estimates for housing prices and transaction volumes, a question remains regarding potential endogeneity concerns inherent in estimating the treatment effects of GSI. Specifically, one might worry that neighborhoods selected for SuDS implementation differ systematically in unobserved ways that also influence housing market trajectories, thus violating the crucial identification assumption underpinning my difference-in-differences approach. To address this concern, I adopt an event study framework, a strategy that helps me test the parallel-trends assumption and explore the dynamic treatment effects over time. I adopt the following event study specification:

$$\ln Y_{i,s,t} = \alpha_s + \lambda_t + \sum_{k=-m}^{-2} \tau_k \Delta D_{s,t}^k + \sum_{k=0}^{q} \tau_k \Delta D_{s,t}^k + \boldsymbol{X}_i^\top \boldsymbol{\gamma} + \epsilon_{i,s,t}. \tag{4}$$

This specification mirrors my baseline two-way fixed-effects model (Eq. (3)), with an important distinction: the treatment indicator is replaced by a sequence of event-time indicators, denoted by the difference operator $\Delta D_{st}^k \equiv D_{s,t-k} - D_{s,t-k-1}$, which equals one if neighborhood (LSOA) s first received SuDS investment exactly k periods before year t. The year immediately preceding treatment ($k = -1$) is omitted as the reference category, normalizing the treatment effect in that year to zero. The estimated coefficients τ_k thus measure the differential trajectory of property values in treated neighborhoods relative to untreated counterparts at various leads ($k < 0$) and lags ($k \geq 0$) around treatment initiation. A key benefit of the event study design lies in its explicit examination of pre-treatment periods, providing a direct assessment of whether treated and untreated neighborhoods exhibited parallel trends prior to the introduction of GSI. If this is the case, it adds up to the credibility of my results.

The event study framework can also inherit limitations of traditional two-way fixed effects estimators under staggered adoption settings (though to a lesser extent), which can inadvertently introduce negative-weighting bias into average treatment effects (Sun and Abraham, 2021). To circumvent this concern, I implement the imputation-based event study method developed by Borusyak et al. (2024), which explicitly estimates counterfactual outcomes for treated neighborhoods using untreated observations and baseline trends, thereby isolating treatment effects from confounding temporal patterns. Specifically, the imputed counterfactual outcome for neighborhood s, treated at event-time k in calendar year t, is constructed as follows:

$$\widehat{Y}_{s,t}^{N(0)} = \hat{\alpha}_s + \hat{\lambda}_t + \boldsymbol{X}_i^\top \hat{\boldsymbol{\gamma}}, \tag{5}$$

[15] The planning permission process itself may become more stringent following SuDS designation. Local authorities, having invested public resources in green infrastructure, face incentives to preserve neighborhood permeability through restrictive zoning that limits density or impervious surface coverage. For instance, evidence from planning applications suggests longer approval timelines and higher rejection rates for developments that might compromise drainage capacity (McLeod and Mickovski, 2024).

where $\widehat{Y}^{N(0)}_{s,t}$ denotes the imputed outcome for treated LSOA s had it never received SuDS treatment by year t, and $\hat{\alpha}_s$, $\hat{\lambda}_t$, and $\hat{\gamma}$ are parameters estimated exclusively from untreated and not-yet-treated neighborhoods. Intuitively, this methodology recovers credible counterfactual trajectories from untreated units, providing a clear baseline against which observed post-treatment outcomes can be compared. Fig. 3 plots the event-time coefficients and their 95 percent confidence intervals from estimating Eq. (4) for the effects of the SuDS program on local housing prices, where period zero denotes the calendar year of the first SuDS delivery within an LSOA. My specification follows column (3) of Table 3, incorporating LSOA and calendar-year fixed effects, year-quarter fixed effects, and the full vector of property-level covariates, with standard errors clustered at the LSOA level.

Fig. 3, Panel A reveals patterns that strongly support my causal interpretation: all pre-treatment coefficients for periods −8 through −1 do not differ significantly from zero. The point estimates in these pre-treatment years range from −0.022 to 0.010, representing economically negligible deviations from the omitted category. A joint F-test of the hypothesis that all pre-treatment coefficients equal zero yields a test statistic of 0.82 ($p = 0.58$), failing to reject the null hypothesis of common pre-trends. The absence of differential pre-treatment trajectories between future-treated and control neighborhoods alleviates the endogeneity concerns that my estimates capture the continuation of pre-existing growth differentials rather than causal infrastructure effects. After the launch of SuDS projects, there is an immediate response, with property values appreciating by around 4.0 percent in the treatment year. The magnitude of the initial effect persists until the third year, when the point estimate decreases and becomes statistically insignificant. Remarkably, treatment effects rise again over the subsequent two years, peaking at 7.8 percent by the fifth year, but estimates then decline from period six onward and exhibit substantially larger standard errors. In Appendix Fig. E.1, I also explore the sensitivity of the results by implementing two additional event study estimators that address the concerns that the linear TWFE regression model may be biased in the presence of staggered treatment timing or treatment heterogeneity: the Sun and Abraham (2021) interaction-weighted estimator and the Callaway and Sant'Anna (2021) aggregation method. The results are not sensitive to the choice of estimators.

This hump-shaped profile is consistent with a forward-looking but imperfect-information view of housing markets rather than a one-time, perfectly permanent capitalization of true risk. In line with recent behavioral models of flood-risk perception (e.g., the myopic learning model in Pryce et al., 2011 and the bounded-Bayesian learning model in Haer et al., 2017), buyers appear to react promptly when visible adaptation works and policy designations make local investment salient, but then continue to revise their beliefs about flood risk and infrastructure performance as they observe how the system behaves in subsequent storms and how reliably it is maintained. What is capitalized into prices at each horizon is thus households' evolving subjective assessment of the expected stream of avoided damages and amenity benefits, conditional on the current configuration of assets and institutions, rather than the elimination of risk once and for all.

In this light, maintenance and institutional support matter not because they overturn the basic capitalization logic, but because they shape beliefs about the persistence of benefits. A credible inspection and maintenance regime, together with stable planning requirements for sustainable drainage, reinforces the perception that SuDS will continue to function over the relevant planning horizon and helps sustain the premium over time. Conversely, if features appear clogged or neglected, or if flood experience contradicts expectations, households may gradually downgrade their beliefs about effectiveness and the premium may attenuate, even though the physical assets remain in place. Later in Section 5.1, I show that capitalization indeed strengthens in neighborhoods that experience the 2021 flash-flood events, consistent with myopic or bounded Bayesian learning in which salient shocks prompt discrete belief revisions rather than once-and-for-all full information.

Next I examine heterogeneity by disaggregating the event study by baseline flood-risk categories (Fig. 3 Panel B). High-risk and low-risk areas display similar trajectories, although the latter shows a larger immediate treatment response and more pronounced decay after five years. The divergence aligns with theoretical predictions in Feng and Nassauer (2022) that functional benefits provide durable value while amenity effects attenuate through habituation or deferred maintenance.[16] In contrast, moderate-risk neighborhoods show no statistically significant effects through year two. Only in year three does a meaningful treatment response emerge, subsequently building to peak effects of around 4.0 percent by year five. Once established, however, these effects prove remarkably persistent through the full post-treatment window without the volatility observed elsewhere. Consistent with the non-monotonic ranking from earlier TWFE estimates in Table 3, these disparate post-treatment trajectories suggest that SuDS capitalization is not simply proportional to actuarial flood risk.

5. Mechanism

Thus far, I have shown that high- and low-flood-risk areas exhibit larger average appreciation effects, whereas moderate-risk areas lag behind. This result, in isolation, is counterintuitive relative to simple "more risk, more payoff" story that underpins most hedonic models of flood exposure. What could possibly generate a U-shaped capitalization schedule? I propose three complementary hypotheses that anchor the subsequent empirical tests.

5.1. Hydraulic-performance heterogeneity

Consider first the functional performance hypothesis, which posits that engineering effectiveness varies systematically across flood risk categories. Under this mechanism, green infrastructure in high-risk and low-risk areas achieves superior hydrological performance compared to installations in moderate-risk neighborhoods. The argument proceeds from recognizing that extreme-risk areas may receive more carefully engineered solutions (e.g. deeper bioretention basins, larger storage volumes, or more permeable surface coverage), precisely because the stakes are higher. On the low-risk margin, many LSOAs in the suburbs might accommodate more ambitious designs therefore generate large, locally salient improvements in drainage. By contrast, moderate-risk tracts in the middle of the distribution often combine modest clay content with legacy gullies that already cope "well enough" in routine storms; incremental SuDS capacity buys little additional protection and, hence, commands only modest willingness to pay.

As a direct test of differential engineering performance, I exploit a natural experiment provided by London's July 2021 extreme precipitation events.[17] I leverage SAR imagery from the ESA's Sentinel-1 satellite constellation to measure road surface water accumulation at 10-meter spatial resolution. The VH (vertical transmit, horizontal

[16] In high-risk areas where annual flooding probability exceeds 3.3 percent, the substantial property value gains clearly reflect capitalization of reduced flood exposure. Yet the parallel trajectory in low-risk neighborhoods — where flooding represents a minimal threat — indicates that environmental amenities, urban cooling, and aesthetic improvements generate comparable willingness to pay.

[17] On July 12 and July 25, 2021, London experienced two severe flash floods in close succession. In some locations, rainfall within a two-hour window exceeded twice the typical monthly total, generating widespread surface-water and sewer inundation. Disruptions included the closure of rail and underground stations and schools, as well as evacuations of hospital wards. PERILS AG subsequently estimated insured losses for the 12–14 July event at approximately £281 million (*Insurance Post*, 13 October 2022, by Pamela Kokoszka).

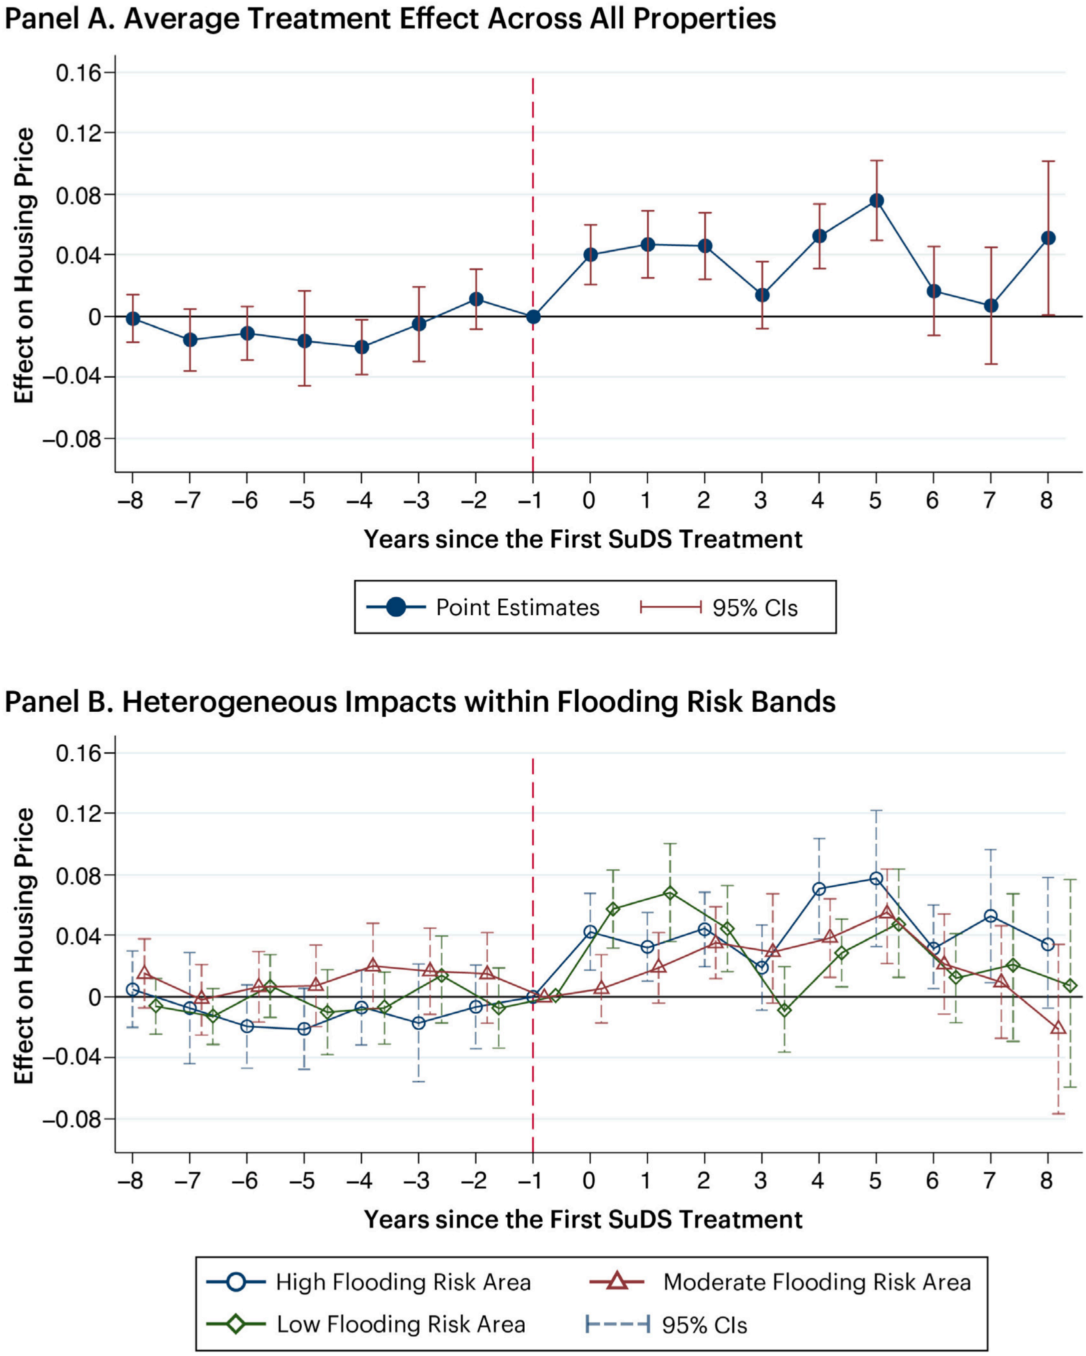

Fig. 3. Dynamic SuDS treatment effects on housing price through event study estimates.
Notes: This figure reports the coefficients of the event time coefficients from event study estimates of log housing transaction price on event-time indicators relative to the year immediately prior to the first SuDS delivery ($k = -1$, omitted). Panel A reports average effects for all treated LSOAs; Panel B estimates the model separately by flood-risk stratum (high, moderate, low) defined using Environment Agency AEP bands. The specification includes LSOA and calendar–year fixed effects, year–quarter fixed effects, and the full set of property-level covariates used in Table 2 Column (3). Standard errors are clustered at the LSOA level. vertical lines denote the treatment year ($k = 0$) with capped bars show 95 percent confidence intervals. Leads are plotted up to eight years before treatment and lags up to eight years after with longer horizons ($k < -8$ and $k > 8$) being binned together. Estimates are obtained via the Borusyak et al. (2024) imputation method, using only untreated observations to fit counterfactual trends.

receive) polarization channel serves as a sensitive indicator of standing water, as smooth water surfaces generate specular reflection that dramatically reduces backscatter returns to the satellite sensor. To quantify infrastructure effectiveness, I calculate the differential flood mitigation during peak inundation periods. Let $\overline{VH}_{i,t}^{pre}$ denote neighborhood i's average dry-baseline value before the storm (July 4–11 and July 20–24), $VH_{i,t}^{peak}$ is its minimum peak value during the events (July 12 and July 25), D_i represents green infrastructure presence before 2021. The peak-event difference-in-differences is estimated as:

$$\delta = E\left[VH_{i,t}^{peak} - \overline{VH}_{i,t}^{pre} | D_i = 1\right] - E\left[VH_{i,t}^{peak} - \overline{VH}_{i,t}^{pre} | D_i = 0\right]. \tag{6}$$

According to the aggregated results reported in Fig. 4 Panel A, during the July 12 event, treated neighborhoods experienced 2.1 dB

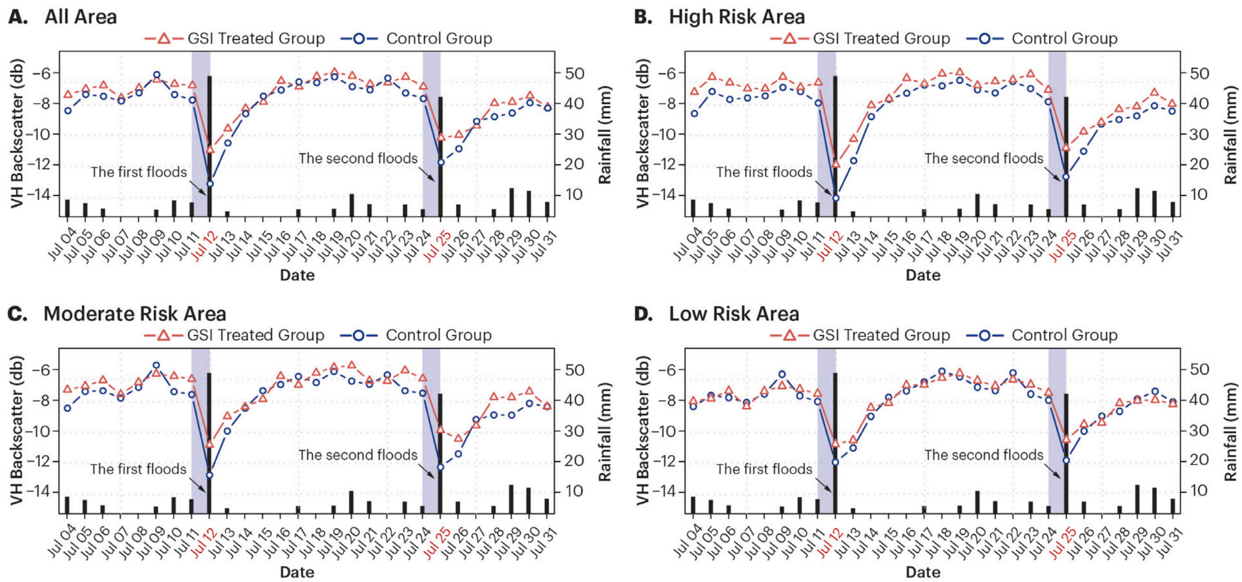

Fig. 4. Road surface water dynamics during two major July 2021 London flash flooding events.
Notes: This figure presents temporal evolution of surface water accumulation measured through Sentinel-1 synthetic aperture radar (SAR) VH backscatter values during two major flooding events on July 12 and July 25, 2021 in Greater London. VH backscatter values (left axis) are plotted as daily averages, with lower values indicating increased surface water presence due to specular reflection. The red triangles represent neighborhoods with GSI installations prior to July 2021, while blue circles show matched control neighborhoods without GSI. The analysis focuses on urban road networks, applying a 10-meter buffer mask around OpenStreetMap road segments to isolate flooding impacts on transportation infrastructure. Daily precipitation totals from the UK Met Office (right axis, black bars). Panel A aggregates all neighborhoods, while Panels B through D stratify by baseline flood risk categories derived from Environment Agency RoFSW data. Shaded regions indicate the flooding event windows. Sample includes 273 treated and 4571 control LSOAs with complete SAR coverage throughout the observation period.

less backscatter reduction than controls, translating to a 23.1 percent larger reduction in roadway surface-water volume. The results confirm that SuDS has meaningful flood-mitigation capacity overall. I then examine the hydraulic-performance hypothesis by comparing δ within each risk band. If superior engineering in extreme-risk areas drove the property-value results, we would expect δ in moderate-risk samples to be significantly smaller than in high- and low-risk categories. However, the estimates prove to be uniform across risk bands: high-risk areas show peak-event protection of 2.3 dB (Panel B), statistically indistinguishable from the 2.2 dB effect in moderate-risk zones (Panel C) and 2.4 dB in low-risk neighborhoods (Panel D). An F-test of equal treatment effects across risk categories yields a statistic of 0.47 ($p = 0.63$), failing to reject the hypothesis of homogeneous infrastructure performance. Therefore, I rule out hydraulic effectiveness as a mechanism explaining non-monotonic capitalizing pattern.

Furthermore, we might want to know whether this realized flood event influences how households capitalize on local adaptation investments. To examine this, I classify neighborhoods as "flood-exposed" or "non-exposed" based on the observed drop in VH backscatter over the road network during the July events, and then re-estimate my housing price regressions both in subsamples and in a pooled specification that interacts SuDS treatment with flood exposure and a post-2021 indicator. I restrict the sample to the 2018–2024 window, and the estimation results are reported in Appendix Table E.7. Consistent with what a myopic risk learning model would predict (Pryce et al., 2011; Lamond et al., 2010), I find little evidence of systematic differences in SuDS capitalization between flood-exposed and non-exposed LSOAs prior to 2021, but a clear divergence opens up thereafter. In the years following this flooding event, the estimated SuDS premium is noticeably larger in LSOAs that experienced visible inundation, whereas changes in non-exposed areas are modest and statistically indistinguishable from the pre-event effects. The evidence suggests that households revise beliefs about flood risk and the value of adaptation most strongly immediately after a salient adverse event.[18]

5.2. Salience and information channel

The second hypothesis shifts attention from hydrology to perception, as different GSI typologies have different visual salience. For instance, a tree-lined bioswale installed along a residential pavement is hard to miss; a sub-surface geocellular tank beneath a car park is invisible. If buyers rely on conspicuous cues when forming beliefs about neighborhood quality, then the market will reward visible GSI adaptations, regardless of their true runoff capture performance. Indeed, buyers can have access to information on SuDS during the purchase process, even though it may not be disclosed to the fullest extent. Since 2016, the standard CON29 property search has included questions about whether a property is served by SuDS, whether features exist within the curtilage, and who bills for surface-water drainage

[18] One caveat concerns potential selective transaction timing following the flood events. If severely damaged properties are more likely to be temporarily withdrawn from the market for repair or insurance settlement, and if damage is correlated with both flood exposure and the absence of effective drainage, then the transacting sample in flood-exposed LSOAs post-2021 may be positively selected toward less-damaged units. Because SuDS placement is itself correlated with baseline surface-water risk, this selection channel is not orthogonal to treatment status and could overstate the estimated post-event SuDS premium in exposed areas.

charges.[19] To bring this mechanism to the data, I classify each installation in the retrofit map into three categories based on salience: (i) *Building-integrated system*: largely roofline or parcel-internal changes with limited streetscape presence. (ii) *Subsurface infiltration infrastructure*: operates below grade or through subtle surface modifications where visual cues are somewhat weak. (iii) *Topographic reconfiguration*: transforms neighborhood landscapes through prominent vegetation, earthworks, and water features that announce their presence to all passersby.

Fig. 5, Panel A, plots the composition of SuDS typologies across baseline flood-risk bands. We find that high- and low-risk areas receive predominantly topographic reconfigurations, with prominent interventions like rain gardens and bioswales comprising over half of installations. On the flip side, moderate-risk neighborhoods disproportionately receive low-visibility systems, concentrated in subsurface and building-integrated systems. One exception is detention basins, for which the share in moderate-risk areas is nearly 60 percent. This sorting pattern aligns with the salience hypothesis: places where I previously measured the strongest capitalization (high- and low-risk) are also the places where residents are most likely to encounter and perceive GSI in everyday experience.

To formally test the information channel and isolate GSI treatment effects by typology, I estimate Eq. (3) in a saturated TWFE model that includes a full set of subtype indicators $D^{g}_{s,t}$ (one for each typology g), together with LSOA and calendar-year fixed effects, year-quarter fixed effects, and the full vector of property covariates. Identification comes from within-LSOA timing as particular subtypes arrive, partialing out the presence of other types. It is noteworthy that the subtype-specific average treatment effects reported in Fig. 5, Panel B, are monotone in typological salience. Interventions that reconfigure the streetscape exhibit the largest and most precisely estimated gains, while building-integrated but invisible systems fail to generate market premiums, with point estimates not significantly different from zero, even though their runoff-capture performance is similar. SuDS projects that are legible to households can raise private willingness to pay. It is also worth noting that effects attributed to salience could, to some extent, reflect amenity improvements, at least for some types of SuDS such as rain gardens. While it is difficult to isolate the salience channel from amenity pathways, the stark heterogeneity across infrastructure types with comparable engineering effectiveness suggests that GSI capitalization is driven less by underlying hydraulic protection but more by what households can perceive, understand, and enjoy in their surroundings. These typological-based results do not imply, however, that invisible GSI is of no value because some long-run socioeconomic benefits may not be reflected in the short-run housing market.

Another concern is that typology choice itself, like the presence of SuDS treatment, may be endogenous to pre-existing neighborhood characteristics — such as density, land use, or land values — that also shape subsequent housing-market trajectories. If, for example, dense commercial LSOAs with limited open space systematically receive subsurface systems, then estimated salience effects might instead be capturing unobserved differences in urban form. To examine the selection mechanism associated with SuDS typology, I fit a set of probit models to investigate the extent to which SuDS typologies are predictable from pre-treatment neighborhood characteristics (see Appendix D for details). Appendix Table D.1 provides compelling evidence that SuDS type assignment is driven primarily by engineering and geomorphological constraints rather than socioeconomic sorting. These findings support the interpretation that visibility itself is an important margin through which green infrastructure is capitalized into housing prices.

5.3. Greenery effects and spatial spillover

Building on the salience and information channel, it is worth asking whether property-value gains may simply reflect urban greening rather than stormwater management per se.[20] Many SuDS introduce substantial vegetation into previously barren streetscapes; if markets primarily value trees, shrubs, and green space rather than hydrological functions, my estimates may conflate distinct amenity channels that deserve separate consideration.[21] To isolate the pure greening channel, I conduct a causal mediation analysis (Imai et al., 2010) by treating greenery as a mediator lying on the causal path between GSI and housing prices and estimating a controlled direct effect (CDE). For each LSOA and year, I use annual Normalized Difference Vegetation Index (NDVI) values $N_{i,t}$ as a proxy for the mediator, using Landsat-8 multispectral data that capture the neighborhood greenness.[22] As a satellite-based index that measures green vegetation, NDVI is an imperfect proxy for *perceived* greenness or design quality, and it may respond to contemporaneous greening initiatives unrelated to SuDS.

Table 4 column (4) presents the net-of-green estimates that decompose the overall SuDS treatment effect into a controlled direct effect and an indirect effect through greenery. The coefficient on treatment indicator $D_{i,t}$ represents the net treatment effect while holding measured greening constant. The greening coefficient, in turn, summarizes the reduced-form price gradient with respect to vegetation (greenery effect). Compared with the baseline estimates in column (1), the greenery channel accounts for approximately one-third of total SuDS value creation. In the full sample (Panel A), controlling for contemporaneous greenery reduces the SuDS coefficient from 0.032 to 0.022 — a 31 percent decline in the treatment effect that can be attributed to greenery channel — with greenery itself contributing 1.2 percent to property values. The sizeable and significant net SuDS effect suggests that stormwater-management benefits extend beyond vegetation enhancement alone, though the decomposition warrants some caution given measurement error or possible post-treatment bias in the NDVI proxy as mentioned before. In panel B to D, I break down the samples by risk stratum. I observe that greenery effects

[19] A CON29 property search is a standard set of enquiries made during a property transaction to uncover crucial information from a local authority, such as details on road proposals, planning history, and building control history. It is part of the official local authority search in the United Kingdom. In practice, however, the information environment is somewhat noisy and incomplete. Because Schedule 3 of the Flood and Water Management Act 2010 has not been implemented in England, local authorities are not legally required to maintain comprehensive SuDS registers, and many explicitly note that their records do not allow them to answer the CON29 questions reliably, advising purchasers instead to check planning approvals and ask vendors directly. As a caveat, information about subsurface systems may be uneven and often requires costly effort to obtain, whereas highly visible interventions automatically signal the presence of adaptation without buyers having to navigate imperfect formal channels.

[20] It is theoretically difficult to separate the greening and stormwater management. As some urban historians pointed out, the emergence of urban greenery, since its beginning of the age of enlightenment, has the dual role of addressing technical challenges and social divisions (Picon, 2014). These two aspects are inextricably linked in forging a political character that exhibits to the present day.

[21] Previous research has identified a positive effect of urban greenery on nearby housing prices. However, the magnitude varies with respect to greenery types and contexts. For example, Trojanek (2016) noted a 3%–4% price increase for homes within 100 m of green spaces in Warsaw, Poland. In the United States, Crompton (2005) found that parks can increase property values by up to 20%

[22] I process all available Landsat 8 scenes from May through September for each year 2010–2024 in London, restricting to images with less than 10 percent cloud cover over Greater London. NDVI values are calculated as $(NIR - Red)/(NIR + Red)$ and aggregated to LSOA boundaries using zonal statistics. Summer observations maximize vegetation vigor and minimize phenological variation.

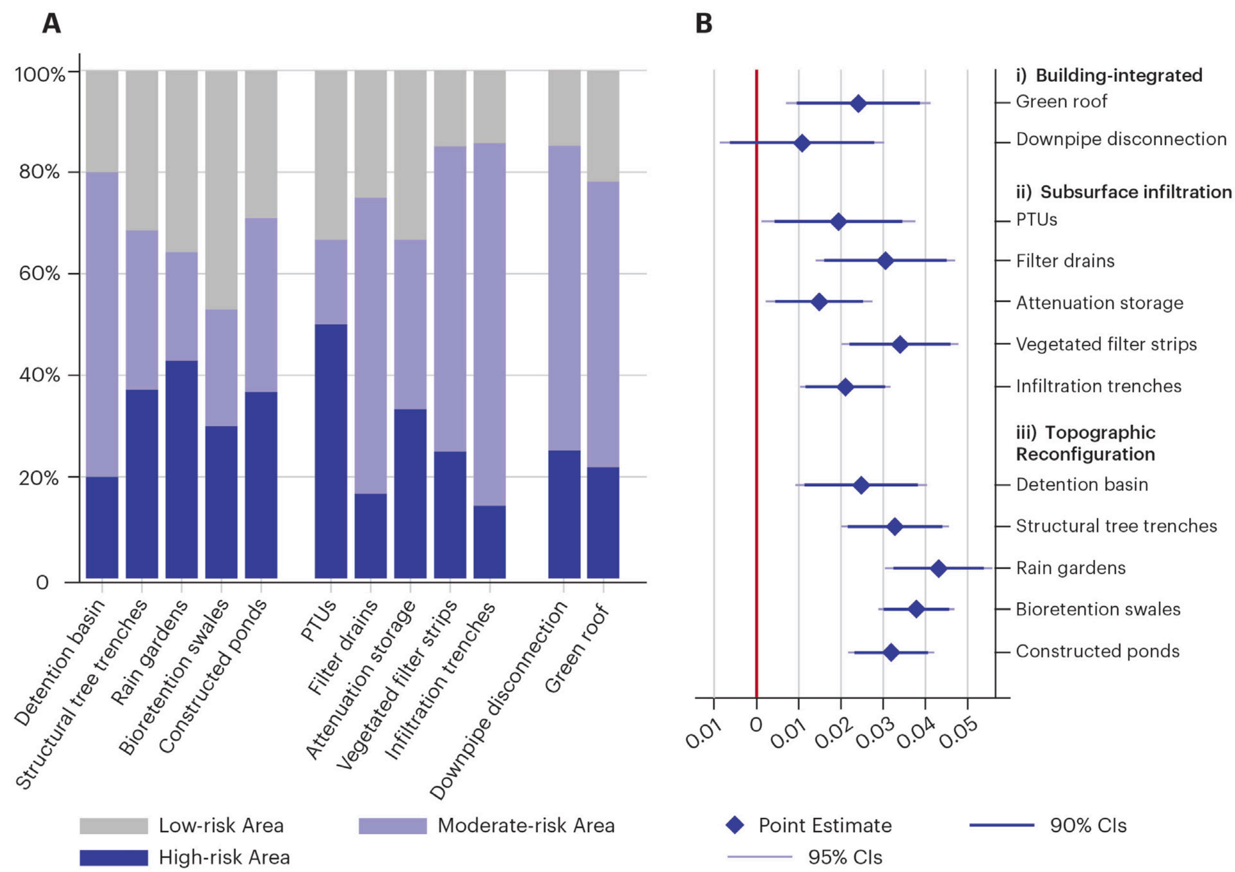

Fig. 5. Composition of SuDS and capitalization effects by typologies.
Notes: Panel A depicts the composition of SuDS installations across baseline flood-risk bands (low, moderate, high) defined by the Environment Agency's AEP thresholds. Each bar reports the share of installations in a typology falling within each risk band. Typologies are grouped by perceptibility. Panel B reports treatment effects on log housing prices from augmented difference-in-differences specifications that interact SuDS implementation with infrastructure type. Point estimates and 90% confidence intervals and 95% confidence intervals are shown for each infrastructure type. Standard errors are clustered at the LSOA level.

are larger for low-risk subgroups, where visible landscape upgrades are most salient. By contrast, in moderate-risk areas, the greenery coefficient is small and fails to achieve statistical significance. This pattern reinforces my earlier findings that moderate-risk neighborhoods receive predominantly invisible infrastructure that neither transforms local greenery nor signals investment to market participants.

Another concern arises from spatial spillovers of GSI hydrological benefits across LSOA boundaries; for example, upstream water capture may protect downstream neighbors regardless of formal treatment status. Additionally, potential slippage in geocoding SuDS installations could misallocate treatments across LSOA boundaries. Both mechanisms would attenuate my estimates by contaminating the control group with partially treated observations. To address these concerns, I exclude all directly treated LSOAs and reclassify their immediate neighbors as "pseudo-treated" based on spatial contiguity rules (*Rook* and *Queen*). Spillover placebo-test results for the baseline and net-of-green specifications appear in Table 4, column (2)–(3) and (5)–(6). These pseudo-treatment effects remain economically small and statistically indistinguishable from zero, with no evidence that spillovers concentrate in areas where hydrological connectivity might be strongest. Hydrological spillovers need not align with administrative adjacency, accordingly, this contiguity-based spillover test is better viewed as a falsification test for highly local contamination (and for geocoding boundary slippage) than as a comprehensive test of basin-level effects.

6. Cost–benefit and public finance perspectives

From the standpoint of public finance, the policymaker's concern extends beyond whether individual property prices rise or fall. The normative question is whether public investment in green stormwater infrastructure generates net social benefits. The SuDS capitalization estimates provide an input into a back-of-the-envelope cost–benefit calculation because, under standard asset-pricing logic, the price change is a stock measure that reflects buyers' discounted expectations about future flows. With approximately 368 treated LSOAs containing an average of 200 residential properties valued at £820,000 each, a 3.5 percent appreciation translates to aggregate value creation of £2.1 billion across treated neighborhoods. This average effect is in fact an indication of a mixture of salient and invisible projects. As the mechanism analysis in Section 5.2 demonstrates, salient streetscape-transforming installations generate capitalization effects three to four times larger than equally effective but inconspicuous alternatives. This heterogeneity implies that municipalities facing tight budget constraints could achieve returns substantially exceeding the average by strategically targeting investments toward visible projects that communicate adaptation effort to prospective buyers.

On the cost side, the funding for London's SuDS program is of a twofold nature, corresponding to the institutional feature as discussed in Section 2.1. The first piece is mandatory, development-driven

Table 4
Net-of-green estimates and neighbor placebo test by flood risk stratum.

	Baseline estimates			Net-of-green estimates		
		Spillover placebo test			Spillover placebo test	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: All LSOA						
SuDS net treatment effects	0.032***	0.005	0.007	0.022***	0.009	0.008
	(0.011)	(0.007)	(0.006)	(0.008)	(0.006)	(0.006)
Greenery indirect effects				0.012**	−0.005	−0.004
				(0.006)	(0.009)	(0.008)
Panel B: High flooding risk LSOA						
SuDS net treatment effects	0.052***	0.009	0.010	0.038**	0.010	0.012
	(0.019)	(0.012)	(0.013)	(0.018)	(0.012)	(0.011)
Greenery effects				0.014**	0.002	0.004
				(0.006)	(0.008)	(0.008)
Panel C: Moderate flooding risk LSOA						
SuDS net treatment effects	0.021*	0.012	0.013	0.017*	0.007	0.005
	(0.012)	(0.009)	(0.008)	(0.010)	(0.009)	(0.008)
Greenery effects				0.004	−0.006	−0.005
				(0.005)	(0.007)	(0.005)
Panel D: Low flooding risk LSOA						
SuDS net treatment effects	0.043***	−0.004	−0.005	0.029***	0.009	0.009
	(0.019)	(0.011)	(0.012)	(0.012)	(0.010)	(0.011)
Greenery effects				0.019***	0.004	0.005
				(0.006)	(0.010)	(0.011)
Spatial contiguity rule	–	Rook	Queen	–	Rook	Queen
LSOA–Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year–Quarter FEs	Yes	Yes	Yes	Yes	Yes	Yes
LSOA linear time trends	No	No	No	No	No	No
Population weights	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each cell reports the TWFE estimates of τ in Eq. (3) using log housing transaction price as dependent variable. Columns (1) present the baseline specification, with the results the same as Table 3 column (3). Columns (2)–(3) are spillover placebo tests that set an LSOA's own treatment to zero and instead assign treatment to adjacent LSOAs based on rook (sharing boundaries) and queen contiguity (sharing boundaries or corners). Columns (4)–(6) augment the baseline models with time-varying Normalized Difference Vegetation Index (NDVI) controls to isolate net treatment effects beyond urban greening. All specifications include LSOA and year fixed effects, year-quarter fixed effects, property-level controls (type, vintage, tenure), and population weights from the 2011 Census. NDVI values are constructed from cloud-screened ($< 10\%$ cover) Landsat 8 imagery during summer months (May-September) for each year 2010–2024. Standard errors clustered at the LSOA level appear in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

provision for new developments since April 2015. Because delivering SuDS on new development, as noted by LSDAP, is often cost-neutral or even cost-saving relative to conventional drainage, the dominant resource cost is most accurately viewed as a private compliance cost borne by developers, as opposed to a cash outflow from borough budgets. The other piece comes from retrofit SuDS, which is financed by multiple funding streams, often part-funded and coordinated through investment programs as detailed further in Appendix Table E.1. Specifically, Department for Environment, Food & Rural Affairs allocated £3.2 million to the GLA for Drain London in 2010, and the Environment Agency reports over £12 million of government investment supporting borough surface-water projects, alongside Thames Water's dedicated "Twenty4twenty" programme (2015–2020) with at least half of its funding deployed in London. These add up to a proximate £30-40 million public investment in total.

To formalize the fiscal implications, consider an idealized *ad valorem* taxation with continuous reassessment. Suppose the estimated aggregate value increment in treated neighborhoods is $\Delta V \approx £2.1$ billion. Applying England's average effective Council Tax rate of $\tau = 0.8\%$, the implied steady-state incremental revenue would be $R = \tau \cdot \Delta V \approx £16.8$ million per year. Against initial capital outlays of £30–40 million, this implies a first-year return of 42–56 percent, substantially exceeding typical public infrastructure hurdle rates. To compute present values, I assume: (i) value appreciation materializes immediately upon SuDS completion and persists without decay; (ii) the tax base adjusts contemporaneously with market values; (iii) a social discount rate of $r = 5\%$ reflecting HM Treasury Green Book guidance; and (iv) a 20-year evaluation horizon corresponding to standard infrastructure appraisal practice. Accordingly, the present value of incremental tax revenues follows the standard annuity formula $PV(\Delta R) = \Delta R \cdot \frac{1-(1+r)^{-T}}{r} \approx 16.8 \times 12.46 \approx £209$ million. Netting the midpoint capital cost of £35 million and assuming annual maintenance expenditures of 2 percent of capital (£0.7 million per year, with present value £8.7 million), the net present value reaches approximately £167 million, implying a benefit–cost ratio of 4.8–6.7 depending on initial outlay assumptions.

This calculation, however, rests on fiscal assumptions that do not hold under current institutions in England. Council Tax, the primary residential property tax, is levied on valuation bands established using 1991 property values. Absent comprehensive revaluation or individual rebanding, market appreciation generates no automatic revenue increase. A property that appreciates from £800,000 to £828,000 following SuDS installation remains in the same Council Tax band, contributing identical revenue to local authorities. The fiscal dividend is therefore unrealized under the existing system. There are nonetheless three alternative ways to capture a portion of the value increment. First, the UK raises substantial revenue through Stamp Duty Land Tax, which is triggered at transaction and can therefore capture part of any SuDS-related appreciation when properties sell (though the amount depends on turnover and the SDLT schedule). Second, the Community Infrastructure Levy (CIL) and Section 106 planning obligations enable local authorities to extract developer contributions for infrastructure provision at the point of development approval. These instruments could, in principle, require developers to fund SuDS as a condition of permission. Third, to the extent that adaptation investments affect commercial property and local economic activity, there may be implications for non-domestic rates (business rates), though the mapping from local value changes to local fiscal capacity is shaped by the details of the retention and equalization regime.

Despite the large benefits of the SuDS program, all is not completely well; there are distributional concerns because costs and benefits do not necessarily accrue to the same individuals. Parallel evidence from the United States by Cohen et al. (2025) analyzes three approaches to financing GSI — flat-rate user fees, proportional property taxes, and impervious surface area charges — and finds all are regressive. This provides a direct caution for interpreting the welfare gains from SuDS. Even if SuDS generate large aggregate benefits, a poorly designed charge can shift costs onto lower-income households, potentially in neighborhoods that may already be prioritized for adaptation. A natural implication is that any discussion of "self-financing" should be paired with explicit design features (i.e., exemptions, income-based rebates, or progressive rate structures) if the policy goal is to expand SuDS while avoiding regressive incidence. More broadly, this distributional caution complements evidence from life-cycle cost analysis that a green/gray infrastructure portfolio can dominate a gray-only benchmark on cost-effectiveness grounds (Cohen et al., 2012). Hence, even when aggregate net benefits are positive, the central finance problem is designing a financing mechanism so that cost recovery does not undo welfare gains through regressive incidence.

7. Conclusion

In the past decade, with the frequency and intensity of adverse climate events escalating, cities have increasingly turned to nature-based stormwater systems to mitigate chronic surface flooding. These interventions are modular and comparatively inexpensive relative to gray infrastructure, but their cumulative footprint and recurrent maintenance needs still require sustained public finance. Understanding how these distributed green infrastructure investments affect local economic activity is critical for governments to justify public expenditure and leverage public–private partnerships to bridge financing gaps. Existing evidence relies on small geographical areas, short-term assessments, and observational hedonic price models whose findings may be substantially driven by selection bias. This paper presents a detailed investigation into London's Sustainable Drainage Systems (SuDS) program, examining how a city-wide public investment in scalable green infrastructure affects local housing markets.

Using a quasi-experimental design, I find evidence that green stormwater infrastructure generates measurable economic returns through property value appreciation. The staggered implementation of SuDS across London's neighborhoods from 2012–2024 increased housing prices by 3.2–3.7 percent on average, with no discernible effects on transaction volumes. These impacts are robust to different specifications and are not affected by the choice of estimator. While treatment assignment is not random, I deploy multiple complementary strategies — combining rich fixed effects, event study estimators, and weighting-based rebalancing — to mitigate observable selection and assess the sensitivity to remaining sources of endogeneity. Therefore, despite some imbalance between treatment and control neighborhoods in the year preceding SuDS rollout with regard to baseline socioeconomic conditions, a causal interpretation of program effects is supported by event study estimates that outcomes in treated and control areas follow parallel trends prior to treatment. Using a spatial placebo test, I further confirm that the results are not affected by slippage of treatment relocation, spillover, or other spatial measurement error.

While municipalities worldwide have committed to increasing green infrastructure provision, budget constraints and uncertainty about economic returns impede implementation at the scale required for climate resilience. My findings provide empirical evidence for policy debates over financing pressing nature-based climate adaptation, but external validity is naturally conditioned by London's distinctive institutional and market environment. Specifically, London combines extreme land scarcity and tight planning constraints, very high incomes, and a housing stock dominated by flats; in addition, flood-insurance institutions (including Flood Re) can mediate how flood risk and risk-reducing investments are priced.[23] Price appreciation without quantity expansion also implies that developers may be less willing to supply new housing due to planning constraints, which may dissipate further value-capture benefits. Future research should examine whether these patterns persist across different institutional contexts, particularly in developing cities with less restrictive planning regimes that might allow supply responses.

I also find evidence of non-monotonic treatment effects across flood risk categories with high-risk and low-risk areas experiencing 4–5 percent gains while moderate-risk neighborhoods show minimal response. My mechanism tests reveal why crude engineering metrics provide poor guidance for infrastructure prioritization. The uniform hydraulic performance across risk categories during the July 2021 flooding events demonstrates that functional effectiveness alone cannot explain heterogeneous market responses. Instead, the dominance of visibility and salience in driving value creation suggests that optimal implementation strategies must balance technical performance with public perception. Highly visible interventions like rain gardens and bioswales generate returns three to four times larger than equally effective but invisible subsurface systems. The lessons have important implications for GSI program design and finance: where budgets are tight, prioritizing salient adaptations and communicating effectively can magnify private willingness to pay. Finally, translating these benefits in a more equitable way and mitigate cost-shifting require careful attention to both market mechanisms and regulatory frameworks.

Declaration of competing interest

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Related work review

This study builds on two complementary strands of empirical literature: the engineering-focused evaluation of green infrastructure performance and the economic assessment of urban climate adaptation investments. Early scholarship on green stormwater infrastructure emerged from hydrological and engineering disciplines, concentrating primarily on technical efficacy (Jayasooriya and Ng, 2014). These studies primarily evaluate the efficacy of green infrastructure by measuring urban stormwater runoff quality and quantity through simulation like catchment modeling tools (Versini et al., 2018; Fowdar et al., 2022).

Most relevant here are studies that move from engineering and environmental metrics to economic valuation. Initial cost–benefit analyses compared green infrastructure investments against conventional gray infrastructure alternatives (Chui et al., 2016; W. Liu et al., 2016), yet these studies struggled to capture the full spectrum of costs and benefits. As Sunstein (2018) notes, traditional hydrologic-focused approaches systematically undervalue non-quantifiable benefits, creating a fundamental challenge for comprehensive economic assessment. This limitation becomes particularly acute given that stormwater infrastructure generates no direct revenues, functioning instead as pure public expenditure requiring justification through indirect benefit channels. Recent analysis from Galvin and BenDor (2023) evaluates the economic impact of green stormwater infrastructure program in Washington, D.C, through job creation and revenue added to the regional economy. Such

[23] These features may shape both the siting and timing of SuDS rollout and the extent to which buyers capitalize future damages and adaptation benefits into transaction prices. Accordingly, the estimated capitalization effects are most likely to generalize to settings where (i) flood risk is salient to households, (ii) adaptation investments are visible and credibly maintained, and (iii) markets are supply-constrained so that amenity and risk reductions are reflected primarily in prices rather than new construction.

analysis still faces an inherent trade-off between comprehensiveness and accuracy when assigning monetary values to non-market amenities (Herath and Bai, 2024).

Two closely related, policy-facing studies illustrate how cost–benefit frameworks have been used to appraise London's SuDS program in practice. Johnson and Geisendorf (2019) develops a neighborhood-level CBA that monetizes a range of ecosystem-service benefits from alternative SuDS portfolios and compares these benefits to costs using net present values and benefit–cost ratios, emphasizing the decision relevance of scenario-based economic appraisal for local stormwater planning. Ossa-Moreno et al. (2017) provides alternative evidence at the catchment scale (Decoy Brook), valuing flood-risk reduction using average annual benefits and appraising "wider benefits" via value transfer. Both studies show that including these co-benefits can materially improve the economic case for SuDS and can help allocate investment across stakeholders in line with the benefits each derives. These CBA exercises, however, necessarily require assigning shadow prices to multiple benefit categories, and the resulting valuations can be sensitive to assumptions and transfer choices, especially for non-market "wider benefits". This motivates complementary *revealed-preference evidence* from the housing market, where capitalization into transaction prices provides a behavior-based valuation of the overall bundle of SuDS benefits.

The empirical evidence examining green infrastructure's place-based economic impacts on property values remains limited and geographically concentrated. Results arising from observational studies is exemplified by Davis et al. (2023), which document a 7 percent general increase in home values associated with constructed wetlands in Avondale, Arizona, with properties directly adjacent to water features commanding 14 percent premiums. Similarly, Park and Kweon (2025) investigates the economic effects of stormwater Best Management Practices (BMPs) on housing sale prices in Washington, D.C., using multiple regression analyses. By contrast, they reported mixed results: proximity to BMPs inside parks increased property values, whereas BMPs placed outside parks or near impervious roads decreased them. Both of these two studies use observational designs that do not fully address potential selection bias arising from non-random site placement and rely heavily on cross-sectional data and lacked longitudinal insights, making it difficult to generalize the findings or attribute causality confidently.

There have been few longitudinal evaluations of GSI programs using panel data. One example is Sohn et al. (2020), who employs spatial hedonic models in a longitudinal framework to evaluate retention pond conversions in Houston, Texas. Their findings of significant positive impacts both cross-sectionally and longitudinally represent an important advance, though the study's focus on low-density suburban subdivisions and single infrastructure type raises questions about generalizability to dense urban contexts with diverse green infrastructure portfolios. The most direct and recent study is that of Cohen et al. (2025), who develops a more convincing quasi-experimental design with DiD estimator for green stormwater infrastructure in Philadelphia. They define the treatment group as properties within 25 m of completed installations and the control group as properties near proposed but not-yet-completed sites. Their estimates indicate substantial price effects, approximately 10 percent for properties in immediate proximity, with clear distance decay to 3.2 percent at 50 m. However, what still remains to be explained is the mechanism through which green infrastructure affects local housing market.

This study advances this literature in several dimensions. First, I exploit quasi-experimental variation from London's staggered SuDS rollout to identify causal effects, addressing the selection bias that limits inference in prior observational studies. Moreover, I complement the canonical difference-in-differences design with coarsened exact matching and inverse probability weighting to address observable selection, providing multiple complementary approaches that yield consistent estimates. Second, unlike prior studies that estimate average treatment effects without investigating underlying channels, I provide direct evidence on mechanisms. Third, I examine heterogeneous effects across flood risk levels, providing insights into the relative importance of functional versus amenity benefits. Finally, my comprehensive dataset spanning over a decade allows examination of both immediate and persistent effects, addressing the temporal limitations of cross-sectional analyses.

Appendix B. Housing price trends by treatment

I motivate the panel identification strategy for the SuDS program with a graphical comparison of housing-price trends in treatment and control samples. The treatment group is those neighborhoods that have implemented SuDS projects by the end of the observed window (2024). The control group consists of areas that have never been exposed to such GSI adaptations. Appendix Fig. B.1, Panel A, shows that housing price in treated samples are slightly higher than in control samples, albeit they move parallel in the pre- SuDS delivery period (no significant differences in pretrends), which is similar to the common trend assumption in DiD assumption. Differences in price growth in the post-treatment period offer compelling evidence of GSI adaptation effect. The gap gradually widens from the first SuDS initiation year through 2019 but shrinks thereafter, likely reflecting staggered adoption and common shocks (e.g. COVID-19). In later estimation, I also allow the SuDS premium to vary across the pre-COVID (before 2020), COVID (2020–2021), and post-COVID (2022 onward) periods in a pooled interaction model. I obtain similar estimates before and after 2020, with a modest attenuation during 2020–2021 consistent with common pandemic shocks (Appendix Table E.3). It should be emphasized that, as a descriptive visualization of price dynamics, this trend plot cannot establish parallel pre-trends; the formal test using an event study framework is provided in Section 4.4. Nonetheless, the patterns suggest that GSI program is associated with stronger housing-price growth than in control areas without such adaptations.

Appendix Fig. B.1, Panel B, documents housing-price trajectories within treatment samples, using event time relative to SuDS treatment for each sample, as is common in the event study literature (Black et al., 2023). Before the SuDS delivery, the trend is flat and indistinguishable from the pre-treatment mean (£0.82 million). Prices begin to appreciate shortly after the intervention, rising by roughly 30 percent above the pre-treatment mean five years after the first SuDS delivery. A close examination of the plot suggests a cyclical pattern after five years; these fluctuations may reflect the fact that initial GSI treatment is not a one-off event but involves ongoing maintenance, as well as subsequent phases of the projects (e.g., infill investments tied to the original scheme). However, this pattern could also reflect macro housing-cycle shocks or other confounding factors. Therefore, simple pre-post treatment comparisons ($\bar{Y}_{Post} - \bar{Y}_{Pre} \approx 0.43$ million pounds) that do not adjust for fixed effects cannot reveal a causal interpretation.

Appendix C. Difference-in-differences with matching and weighting estimators

The validity of the identification strategy rests on parallel pre-trend which are most clearly assessed in Fig. 3A. However, Table 1 shows that treated LSOAa are younger, more ethnically diverse, more public-transport oriented, and socioeconomically disadvantaged at baseline. That may still raise concerns on unconfoundedness claim. Arguably, My rich set of covariates allows me to bolster the DiD design with matching and weighting approaches that address the most obvious confounders. In particular, one might generate a more credible comparison group, via matching or weighting, that is similar to the treated group that were observationally similar pre-treatment.

I first choose Coarsened Exact Matching (CEM) to construct such a comparison group, because it directly enforces overlap and balance on the specific baseline characteristics that likely governed program

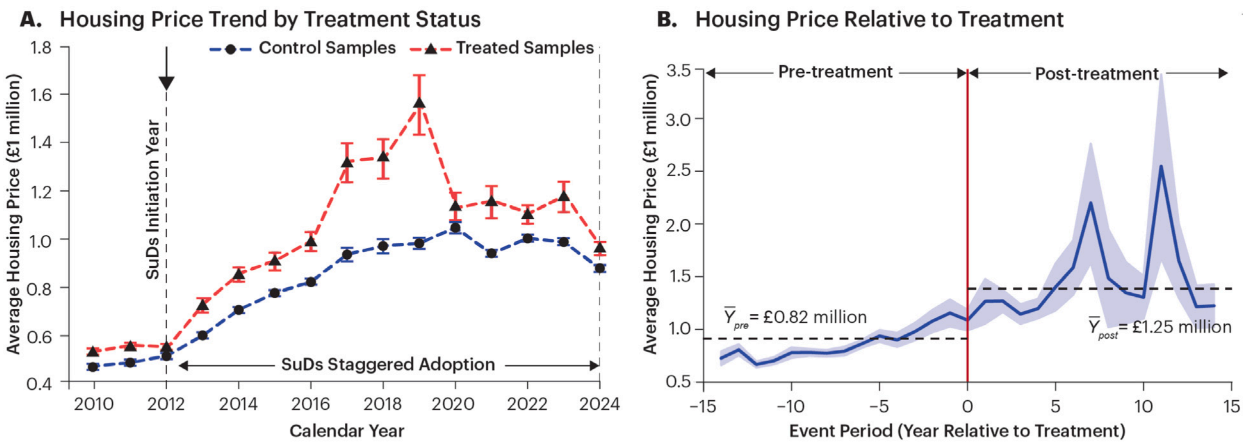

Fig. B.1. Trends in Average housing price at the LSOA level by treatment status 2010–2024.
Notes: This figure presents trends in average housing prices across treated and control samples (Panel A) and relative to the timing of treatment (Panel B). In Panel A, treated samples (N=113,857) include properties in LSOAs that received SuDS adaptation during the study period, while control samples (N=867,405) refer to properties in LSOAs that never received such interventions. Prices are normalized and displayed in millions GBP. Error bars denote ±1.96 standard errors of the mean. Panel B restricts to treated samples and plots housing prices by event time, defined as years relative to the first year of treatment. The shaded region indicates the 95 percent confidence interval. Vertical and horizontal dashed lines mark the timing of treatment and mean housing price, respectively.

targeting. Let D_i in 0, 1 indicate whether an LSOA ever receives a SuDS intervention over the study window, and let $\boldsymbol{X}_i$ denote the vector of baseline (2011) characteristics assembled from the Census and transport records. CEM proceeds by discretizing each continuous covariate into substantively meaningful bins and then forming strata by the Cartesian product of these bins across covariates. Within each stratum, matching is exact on the coarsened variables. Strata that contain only treated or only controls are dropped (non-overlap is pruned), yielding a set of overlap strata in which treated and controls coexist.

Formally, for each covariate X_{sj} I define a binning map $c_j(\cdot)$ that partitions its support into B_j bins; the stratum for LSOA s is $g(s) = (c_1(X_{s1}), \dots, c_J(X_{sJ}))$. Let $\mathcal{G}^*$ denote the subset of strata with at least one treated and one control LSOA. For each $g \in \mathcal{G}^*$, let n_{1g} and n_{0g} be the numbers of treated and control units, respectively. CEM induces simple, transparent weights that reproduce the treated covariate distribution within each overlap stratum. I keep all treated LSOAs in overlap strata with unit weight, assign each control LSOA in stratum g the weight n_{1g}/n_{0g}, and assign weight zero to units in non-overlap strata. Intuitively, within each stratum, the weighted number of controls equals the number of treated; across strata, the overall weighted control distribution mirrors the treated distribution on the coarsened covariates. When I move to micro-data regressions, the CEM weight for LSOA s is attached to all observations from that LSOA over time.

The resulting analysis targets an average treatment effect on the treated (ATT) for the subset of treated LSOAs that have comparable controls in 2011 (the matched common-support region). Identification requires a stratum-specific common-trends assumption: in the absence of treatment, the expected change in outcomes between adjacent pre-periods would have been the same for treated and control LSOAs within each stratum. For all $g \in \mathcal{G}^*$ and pre-treatment t, the assumption is given as:

$$E[\,Y_{s,t}(0) - Y_{s,t-1}(0) \mid D_i = 1,\ g(s) = g\,] = E[\,Y_{s,t}(0) - Y_{s,t-1}(0) \mid D_i = 0,\ g(s) = g\,]. \tag{C.1}$$

Note that CEM does not guarantee this condition, but by enforcing close comparability on covariates that plausibly drove selection to limited extent, it makes the assumption more credible and circumscribes extrapolation.

Finally, As a complementary design-based robustness check, I implement an inverse probability weighting (IPW) estimator that reweights controls to reproduce the baseline covariate distribution of treated LSOAs. More precisely, I first estimate a propensity score $\hat{\pi}(\boldsymbol{X}_s) = \Pr(T = 1 \mid \boldsymbol{X}_s)$ for each observation from a logistic regression that include all pre-treatment covariates linearly. The additional assumption for using propensity score is overlap ($0 < \pi(\boldsymbol{X}_s) < 1$ for all relevant $\boldsymbol{X}_s$). I checked this assumption in Appendix Fig. C.1 which suggests common support: for each level of $\boldsymbol{X}_s$ there is a positive probability of being assigned to treatment. Unlike matching, IPW retains (nearly) the full sample while addressing selection on observables by construction and aggregates smoothly across many covariates. To mitigate the instability of Horvitz–Thompson weights, I use the Hájek normalized form following L. Liu et al. (2016) throughout and integrate these weights into Eq. (3) used in the main text. Both the matching and weighting techniques pursue the same end (constructing credible counterfactuals for treated LSOAs), but through different means. I report both their results to show robustness of the estimation and strengthen the paper's identification argument.

Appendix Fig. C.2 presents measures of covariate balance between GSI treated and control LSOAs before and after matching and weighting procedures. For each covariate, I plot the standardized bias as measured by difference in means (treated minus control) scaled by pooled standard deviation. Vertical dashed lines mark ±0.10 as a conventional balance benchmark; the solid line at 0 indicates exact balance. As expected, there are sizeable baseline imbalances (indicated by gray line and circles) in the distribution of preexisting characteristics between treated and control groups. After CEM, the sample achieves a high degree of balance: all covariates lie within ±0.10, and most are within ±0.05. The largest pre-match imbalances (e.g., car ownership, poor access, social-rented) are reduced to small residuals near zero. IPW (Hájek) delivers a similar balance profile, except for two demographics (i.e., URM, unemployment) falling slightly outside the 0.1 threshold. Differences between the two designs are minor relative to the improvement over the unmatched sample.

Results in Appendix Table C.1 show that estimates from baseline DiD, CEM-then-DiD, and IPW-DiD are remarkably consistent in terms of the sign, magnitude, and risk-layer pattern. SuDS adoption is associated with roughly 3–5 percent higher transaction prices, with the

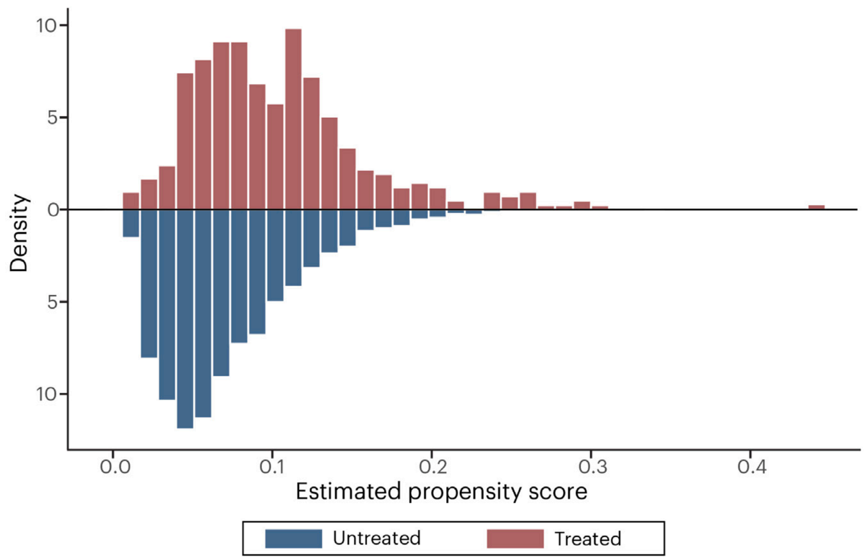

Fig. C.1. Common support of propensity scores.
Notes: This figure plots mirrored histograms of the estimated propensity score $\hat{\pi}(\boldsymbol{X}_i)$, obtained from a logit model using 2011 pre-treatment covariates. Bars above the zero line show the treated distribution; bars below the line show the control distribution. Densities are normalized within group and the buns are 40 equal-width. Overlap is concentrated roughly between 0.03 and 0.20, with a small treated right tail; in the weighting analysis I use Hájek-normalized ATT weights and report robustness to trimming of extreme scores.

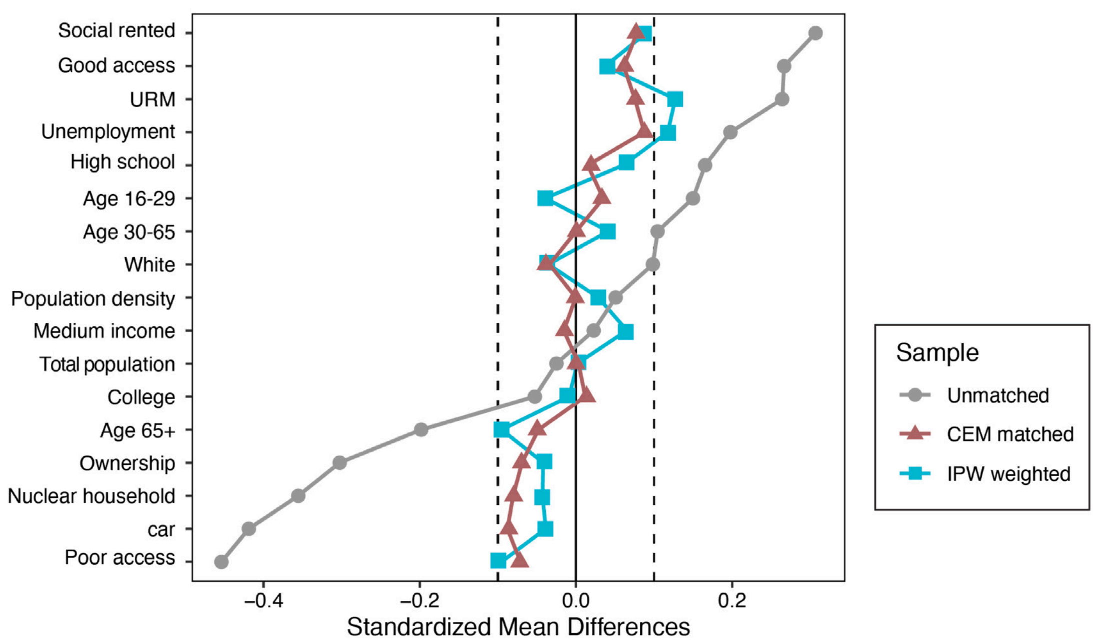

Fig. C.2. Covariate balance before and after matching and weighting design.
Notes: The plot reports signed standardized mean differences (treated minus control, pooled-SD denominator) for 2011 baseline covariates. Gray circles show the raw (unmatched) sample; red triangles show the CEM-matched sample obtained by coarsening all baseline covariates and pruning non-overlap strata, with controls weighted to match treated counts within strata; blue squares show the IPW (Hájek-normalized ATT) sample using a logit of ever-treated on all baseline covariates. Vertical dashed lines mark ±0.10 as a conventional balance benchmark; the solid line at 0 indicates exact balance. Positive values imply higher means among treated LSOAs. After either design adjustment, the majority of covariates fall within ±0.10 and most within ±0.05.

largest gains in low-risk and high-risk LSOAs and smaller, borderline-significant gains in moderate-risk areas. Because CEM hard-trims off-support comparisons and IPW continuously reweights controls to the treated covariate distribution, concordant estimates across these very different design logics substantially mitigate concerns that the main results reflect selection on observables or extrapolation outside common support.

Appendix D. Baseline predictors of SuDS typologies

A central component of my mechanism analysis is the claim that the capitalization of green stormwater infrastructure depends on the salience of SuDS typologies. Fig. 5 shows that high- and low-risk areas are more likely to receive visually prominent "topographic reconfiguration" interventions, while moderate-risk neighborhoods disproportionately receive less visible subsurface or building-integrated systems.

Table C.1
Robustness to design adjustments: coarsened exact matching and inverse-probability weighted DiD.

Risk layer	(1) All	(2) High-risk	(3) Moderate-risk	(4) Low-risk
Panel A. Baseline staggered DiD estimations				
DiD estimator $\hat{\tau}$	0.038***	0.043***	0.016*	0.046***
	(0.008)	(0.016)	(0.009)	(0.012)
Pre-program mean (logarithmic price)	14.021	13.885	14.120	14.003
Observations (N)	981,250	305,003	496,515	179,732
Panel B. Coarsened exact matching then DiD estimation				
DiD estimator $\hat{\tau}$	0.031***	0.041***	0.020	0.042***
	(0.012)	(0.015)	(0.014)	(0.018)
Pre-program mean (logarithmic price)	14.038	13.782	14.201	13.988
Observations (N)	251,563	94,478	39,951	39,951
Panel C. Inverse probability weighted DiD estimation				
DiD estimator $\hat{\tau}$	0.035***	0.039***	0.024*	0.051***
	(0.007)	(0.012)	(0.013)	(0.014)
Pre-program mean (logarithmic price)	14.021	13.885	14.120	14.003
Observations (N)	981,250	305,003	496,515	179,732
LSOA–Year fixed effects	Yes	Yes	Yes	Yes
Year–Quarter fixed effects	Yes	Yes	Yes	Yes
LSOA linear time trends	No	No	No	No
Population weights	No	No	No	No

Notes: Each cell reports the coefficient $\hat{\tau}$ from a staggered difference-in-differences regression of SuDS treatment on log transaction price. Panel A reproduces my preferred baseline specification as in Table 3 column (2). Panel B implements CEM on 2011 baseline demographics, tenure, education, unemployment, income, population density, cars per household, and PTAL. Non-overlap strata are pruned and controls are reweighted within strata (ATT target) so that DiD estimator is re-estimated on the matched common-support sample using CEM weights. Panel C implements IPW (Hájek) using a logit propensity of ever-treated on the same baseline covariates; treated units receive weight 1 and controls receive $\hat{\pi}/(1-\hat{\pi})$, normalized within group; the weighted DiD is estimated on the full panel. Pre-program mean is the average log price in pre-treatment periods within each estimation sample. Risk layers (High, Moderate, Low) follow the Environment Agency's AEP bands. All models include LSOA fixed effects and calendar-time fixed effects (year and year-by-quarter as indicated); standard errors are clustered at the LSOA level. For interpretability, τ corresponds to $100 \times [\exp(\tau)-1]$ percent.

A key concern is that the type of SuDS is not random but is itself a function of neighborhood characteristics that are likely correlated with housing price trends, such that the salience effects could be conflated by the effect of these unobserved characteristics.

To address this concern, I perform a prediction test to check whether the allocation of visible versus invisible SuDS types is tightly linked to baseline observables that are themselves plausible drivers of housing price dynamics. If typology is only weakly correlated with such characteristics, this would mitigate worries that the salience mechanism is merely relabeling selection on neighborhood fundamentals. To do so, I estimate single-equation probit models at the LSOA level. For each typology $g \in \{\text{building-integrated}, \text{subsurface}, \text{topographic}\}$, I define a binary indicator $Y_{sg}=1$ if LSOA s ever receives at least one installation of type g between 2012–2024, and $Y_{sg}=0$ otherwise, restricting the sample to LSOAs that receive at least one SuDS of any kind. I model the selection of typology g using a latent-index framework:

$$Y^*_{sg} = \boldsymbol{X}_s^\top \boldsymbol{\beta}_g + \boldsymbol{B}_s^\top \boldsymbol{\delta}_g + \epsilon_{sg}, \tag{D.1}$$

where $\boldsymbol{X}_s$ is a vector of pre-treatment characteristics, $\boldsymbol{B}_s$ is a vector of borough fixed effects, and ϵ_{sg} is a standard normal error. The components of $\boldsymbol{X}_s$ include: demographic structure (population, density, race and age shares, nuclear household share), socioeconomic measures (homeownership, social renting, unemployment, education, median income), transport and accessibility (cars per household, PTAL 0–1 and 4–6 indicators), urban form and land cover (shares of residential, commercial/industrial, road/transport, open space, and impervious surface),[24] as well as pre-2010 log mean housing price (period 2008–2010).

[24] Land-use shares are constructed from the Ordnance Survey MasterMap Topography Layer, aggregating parcel-level land-use classes to LSOA boundaries and computing the fraction of total land area in each category. Impervious surface share is derived from the Copernicus High Resolution Layer "Imperviousness" 2009 raster product at 10 m resolution, re-sampled and averaged to LSOA polygons.

The latent variable Y^*_{sg} can be interpreted as the net benefit to planners of deploying typology g in LSOA s, relative to not using that typology. Under the normality assumption on ϵ_{sg}, the probability that LSOA s ever receives typology g conditional on its baseline characteristics is given as:

$$\Pr\left(Y_{sg}=1 \mid \boldsymbol{X}_s, \boldsymbol{B}_s\right) = \Pr\left(Y^*_{sg} > 0 \mid \boldsymbol{X}_s, \boldsymbol{B}_s\right) = \Phi\left(\boldsymbol{X}_s^\top \boldsymbol{\beta}_g + \boldsymbol{B}_s^\top \boldsymbol{\delta}_g\right), \tag{D.2}$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function. Parameters $\boldsymbol{\beta}_g$ and $\boldsymbol{\delta}_g$ are estimated by maximum likelihood. As my primary goal is to assess how strongly typology choice is related to observables rather than providing a structural model of the planning decision, I summarize the results using average marginal effects. Specifically, for each continuous covariate x_{ks} in X_s, the marginal effect on the probability of receiving typology g is calculated as:

$$\frac{\partial \Pr(Y_{sg}=1 \mid \boldsymbol{X}_s, \boldsymbol{B}_s)}{\partial x_{ks}} = \phi\left(\boldsymbol{X}_s^\top \boldsymbol{\beta}_g + \boldsymbol{B}_s^\top \boldsymbol{\delta}_g\right)\beta_{kg}, \tag{D.3}$$

where $\phi(\cdot)$ is the standard normal density function. I report the sample average of this derivative across all treated LSOAs as the average marginal effect for covariate k.

Appendix Table D.1 shows the estimates of probabilities and associated marginal effects from the probit model. Across all three SuDS types, demographic and socio-economic variables in general play a limited role. Coefficients on race shares, age structure, tenure, unemployment, education, and income are uniformly small in magnitude and statistically indistinguishable from zero. The few marginally significant coefficients are from population density, which suggest that topographic reconfiguration may prefer low-density neighborhoods to some extent. Moreover, baseline housing market conditions are not able to predict typology as pre-2010 mean house prices have negligible and insignificant effects on the probability of receiving any SuDS type.

Land-use and engineering-related variables, on the contrary, are correlated with typology choice (except for building-integrated systems), in line with the program design conditioned on engineering constraints and hydrological requirements. LSOAs with higher shares of road and transport land and higher impervious surface shares are substantially

Table D.1
Baseline predictors of SuDS typologies at the LSOA level.

	Topographic reconfiguration		Subsurface infiltration		Building-integrated	
	Probit coefficient (1)	Marginal effect (2)	Probit coefficient (3)	Marginal effect (4)	Probit coefficient (5)	Marginal effect (6)
A. Demographic and socio-economic						
Population density (per ha)	−0.021*	−0.008*	−0.015	−0.006	0.009	0.004
	(0.012)	(0.005)	(0.009)	(0.004)	(0.011)	(0.005)
Share White	0.048	0.019	−0.062	−0.024	0.033	0.013
	(0.095)	(0.037)	(0.088)	(0.034)	(0.090)	(0.035)
Share URM	0.083	0.032	0.057	0.022	0.069	0.027
	(0.098)	(0.038)	(0.091)	(0.035)	(0.094)	(0.037)
Share age 16–29	0.061	0.024	0.044	0.017	0.018	0.007
	(0.112)	(0.044)	(0.104)	(0.040)	(0.107)	(0.041)
Share age 65+	−0.057	−0.022	0.039	0.015	−0.028	−0.011
	(0.101)	(0.039)	(0.095)	(0.037)	(0.097)	(0.038)
Homeownership rate	−0.034	−0.013	0.026	0.010	0.019	0.008
	(0.077)	(0.030)	(0.071)	(0.028)	(0.073)	(0.029)
Social rented share	0.069	0.027	0.041	0.016	0.058	0.023
	(0.081)	(0.032)	(0.075)	(0.029)	(0.078)	(0.031)
Unemployment rate	0.118	0.046	0.072	0.028	0.051	0.020
	(0.093)	(0.036)	(0.087)	(0.033)	(0.089)	(0.034)
Qualification ≤ high school	0.064	0.025	0.049	0.019	0.031	0.012
	(0.082)	(0.032)	(0.076)	(0.030)	(0.078)	(0.030)
Qualification = university	−0.079	−0.031	−0.054	−0.021	−0.046	−0.018
	(0.086)	(0.033)	(0.080)	(0.031)	(0.082)	(0.032)
Median household income	0.014	0.006	0.009	0.004	−0.011	−0.004
	(0.011)	(0.009)	(0.012)	(0.006)	(0.010)	(0.006)
B. Land use and imperviousness						
Residential land share	0.176	0.069	0.093	0.036	0.058	0.023
	(0.119)	(0.047)	(0.111)	(0.043)	(0.113)	(0.044)
Commercial land share	0.142	0.056	0.191*	0.074*	0.087	0.034
	(0.117)	(0.046)	(0.111)	(0.042)	(0.113)	(0.044)
Road/transport land share	0.213*	0.083*	0.301**	0.117**	0.094	0.037
	(0.121)	(0.047)	(0.120)	(0.046)	(0.118)	(0.046)
Open space share	0.284**	0.110**	0.097	0.037	−0.062	−0.024
	(0.124)	(0.048)	(0.115)	(0.044)	(0.118)	(0.045)
Impervious surface share	0.421***	0.165***	0.353**	0.138**	0.104	0.041
	(0.140)	(0.055)	(0.130)	(0.051)	(0.134)	(0.053)
C. Transportation and accessibility						
Cars per household	−0.038	−0.015	0.027	0.011	0.019	0.008
	(0.071)	(0.028)	(0.067)	(0.026)	(0.069)	(0.027)
PTAL = 0–1 (poor access)	−0.066	−0.026	−0.041	−0.016	0.058	0.023
	(0.095)	(0.037)	(0.089)	(0.034)	(0.092)	(0.035)
PTAL = 4–6 (good access)	0.054	0.021	0.061	0.024	0.083	0.032
	(0.089)	(0.035)	(0.084)	(0.033)	(0.086)	(0.034)
D. Baseline housing market						
Pre-2010 mean housing price	0.017	0.007	0.012	0.005	0.009	0.004
	(0.041)	(0.016)	(0.038)	(0.015)	(0.039)	(0.015)
Borough fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Log-likelihood	−154.2		−161.8		−147.9	
Observations (LSOA)	368	368	368	368	368	368

Notes: Each column reports coefficients from a separate probit regression at the LSOA level. The dependent variable equals one if an LSOA ever receives at least one SuDS installation of the indicated typology between 2012 and 2024, and zero otherwise; the sample is restricted to LSOAs that receive at least one SuDS of any type. All covariates are measured at baseline (2011 Census and pre-2010 period). Land-use shares are computed from Ordnance Survey MasterMap and the London Development Database; impervious surface share is derived from Copernicus Imperviousness Layer data. Median household income is expressed in units of €10,000. Reported marginal effects are evaluated at the sample means. Borough fixed effects are included in all specifications. Robust standard errors are shown in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

more likely to receive both topographic reconfigurations and subsurface infiltration systems; a 10 percentage-point increase in impervious cover, for example, is associated with roughly a 1.6 percentage-point increase in the probability of a topographic installation and a 1.4 percentage-point increase for subsurface systems. Road and transport land shares similarly predict subsurface systems (marginal effect = 0.117, $p < 0.05$), suggesting subsurface infiltration being the primary feasible option beneath paved infrastructure. The open space coefficient for topographic reconfiguration warrants careful interpretation. While open space availability could proxy for neighborhood amenities that influence property values, its effect likely operates through an engineering channel because pre-treatment housing prices show no significant relationship with SuDS typology selection.

Appendix E. Additional figures and tables

See Tables E.1–E.7 and Fig. E.1.

Data availability

Data will be made available on request.

Table E.1
Funding architecture for SuDS program in London.

Source	Allocation method	Cost coverage	Government level
Channel A. New Development-Driven SuDS			
Planning requirements	Provided as a condition of major development approval; SuDS strategy/pro forma submitted with planning application; LLFA consulted through the planning process	Typically 100% of delivery cost borne by developer (compliance cost; not reimbursed as a grant)	Local (borough LLFA)
CIL/Section 106	Negotiated/levied through development approval; financial contributions may be directed to off-site mitigation measures	Project-specific; partial (varies by agreement/charge schedule)	Local (borough LLFA)
Channel B. Retrofit SuDS			
Drain London	Awarded through a London-wide bid/programme grant; used for mapping, planning, and demonstration projects	Programme grant; not a universal per-project reimbursement (amount varies by programme scope)	Central (Defra/GLA)
FDGiA	Project-based applications/business cases assessed under national FCERM appraisal and partnership-funding rules	Partial to substantial; depends on benefit–cost case and partnership funding (often requires contributions/match)	Central (DEFRA via EA)
RFCC local levy	Committee allocates levy-funded capital to selected flood risk management projects (often used to fill funding gaps)	Partial or full for eligible projects; amounts vary by region and project	Regional
Thames Water	Time-bounded calls for projects; applications assessed and scored within programme rounds	Up to 100% of eligible costs in some cases; frequently co-financed alongside other sources	Non-government
DfE Schools	Competitive, multi-stage grant process with eligibility and award criteria	Up to 50% of total project cost; remaining share secured as match funding	Central

Notes: The table summarizes institutional channels for SuDS delivery and retrofit financing in London. New-development SuDS are primarily delivered through the planning system as a developer-funded requirement, whereas retrofit SuDS in the existing public realm draw on multiple project-based funding streams that commonly involve appraisal and co-financing; cost coverage varies with programme rules and project characteristics. Acronyms: SuDS = Sustainable Drainage Systems; CIL = Community Infrastructure Levy; Section 106 = planning obligations under Section 106 of the Town and Country Planning Act 1990; LLFA = Lead Local Flood Authority; Drain London = London-wide surface water management programme led by the Greater London Authority (GLA) with Defra support; FDGiA = Flood Defence Grant -in-Aid; FCERM = Flood and Coastal Erosion Risk Management; RFCC = Regional Flood and Coastal Committee; Defra = Department for Environment, Food & Rural Affairs; GLA = Greater London Authority; EA = Environment Agency; DfE = Department for Education.

Table E.2
Heterogeneity of SuDS capitalization effects by transaction type (placebo/expected magnitudes).

	Split-sample estimates		Pooled interaction model	
	Standard (Cat. A) (1)	Non-standard (2)	All transactions (3)	High-risk areas (4)
Panel A. Dependent variable: log transaction price				
SuDS × Post (β_1)	0.038***	0.013	0.039***	0.044***
	(0.008)	(0.010)	(0.008)	(0.016)
SuDS × Post × Non-standard (β_2)			−0.026**	−0.033*
			(0.011)	(0.020)
Implied SuDS effect for non-standard ($\beta_1 + \beta_2$)			0.013	0.011
			(0.010)	(0.015)
Observations	981,250	165,420	1146,670	360,180
Panel B. Dependent variable: annual transactions (PPML), by transaction type				
SuDS × Post (β_1)	0.002	−0.021	0.004	−0.070
	(0.036)	(0.052)	(0.036)	(0.045)
SuDS × Post × Non-standard (β_2)			−0.018	−0.041
			(0.051)	(0.073)
Implied SuDS effect for non-standard ($\beta_1 + \beta_2$)			−0.014	−0.111
			(0.049)	(0.069)
Observations	165,167	165,167	165,167	51,514
LSOA–Year fixed effects	Yes	Yes	Yes	Yes
Year–Quarter fixed effects	Yes	Yes	Yes	Yes
Property-level controls	Yes	Yes	Yes	Yes

Notes: "Standard" transactions correspond to arm's-length market sales (Category A). "Non-standard" transactions include Category B and other flagged non-arm's-length transfers (e.g., repossessions, transfers to non-private entities, and investor/buy-to-let purchases where observed). Columns (1)–(2) estimate the baseline specification separately by transaction type. Columns (3)–(4) estimate a pooled interaction model; β_1 is the SuDS effect for standard transactions (omitted group), and the implied effect for non-standard transactions equals $\beta_1 + \beta_2$. Panel B repeats the exercise for market activity using PPML. Standard errors are clustered at the LSOA level and reported in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table E.3
SuDS capitalization before, during, and after the COVID-19 period.

	Dependent variable: log housing transaction price			
	All LSOA (1)	High risk (2)	Moderate risk (3)	Low risk (4)
$D_{s,t}$ × $1[t \leq 2019]$	0.037***	0.041**	0.017*	0.047***
	(0.009)	(0.018)	(0.010)	(0.015)
$D_{s,t}$ × $1[2020 \leq t \leq 2021]$	0.018*	0.019	0.010	0.021
	(0.011)	(0.022)	(0.013)	(0.019)
$D_{s,t}$ × $1[t \geq 2022]$	0.040**	0.046**	0.019	0.041**
	(0.017)	(0.020)	(0.014)	(0.018)
Observations	981,250	305,003	496,515	179,732
LSOA–Year fixed effects	Yes	Yes	Yes	Yes
Year–Quarter fixed effects	Yes	Yes	Yes	Yes
Property-level controls	Yes	Yes	Yes	Yes

Notes: This table reports two-way fixed-effects estimates of how the capitalization of SuDS varies across the pre-COVID, COVID, and post-COVID periods. The dependent variable is the natural logarithm residential transaction prices. The SuDS treatment indicator $D_{s,t}$ equals one from the first calendar year in which any SuDS installation is completed within LSOA s and remains one thereafter. I allow the treatment effect to vary across three time windows by interacting $D_{s,t}$ with indicators for pre-2019 transactions ($1[t \leq 2019]$), the COVID period ($1[2020 \leq t \leq 2021]$), and the post-COVID period ($1[t \geq 2022]$). Coefficients in the first row capture the average SuDS premium in the pre-COVID period; the second and third rows report the corresponding effects during and after COVID-19. All specifications include LSOA fixed effects, calendar-year (or year–quarter) fixed effects, and the full set of property-level controls (type, vintage, tenure). Standard errors, clustered at the LSOA level, are reported in parentheses. Coefficients may be converted to percentage changes as $100 \times [\exp(\beta) - 1]$. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table E.4
Heterogeneity of SuDS capitalization effects by property type.

	Split-sample estimates		Pooled interaction model	
	Houses (1)	Flats (2)	All properties (3)	High-risk properties (4)
SuDS × Post (β_1)	0.045***	0.016**	0.046***	0.054***
	(0.010)	(0.008)	(0.010)	(0.018)
SuDS × Post × Flat (β_2)			−0.028***	−0.029*
			(0.009)	(0.017)
Implied SuDS effect for flats ($\beta_1 + \beta_2$)			0.018**	0.025*
			(0.008)	(0.014)
Observations	342,771	638,455	981,250	305,003
LSOA–Year fixed effects	Yes	Yes	Yes	Yes
Year–Quarter fixed effects	Yes	Yes	Yes	Yes
Property-level controls	Yes	Yes	Yes	Yes

Notes. The dependent variable is the natural logarithm of transaction price. Columns (1)–(2) estimate the baseline SuDS two-way fixed-effects specification separately for houses (including detached, semi-detached, and terraced housing) and flats. Columns (3)–(4) estimate a pooled model that interacts the treatment indicator with an indicator for flats; β_1 is the SuDS effect for houses (omitted group), and the implied SuDS effect for flats equals $\beta_1 + \beta_2$. Column (4) restricts the sample to LSOAs classified as high flooding risk using the Environment Agency AEP bands, consistent with Table 3. Standard errors are clustered at the LSOA level and reported in parentheses. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table E.5
Heterogeneity in SuDS capitalization before and after Flood Re.

	All LSOA (1)	High risk (2)	Moderate risk (3)	Low risk (4)
SuDS treatment $D_{s,t}$	0.034***	0.042**	0.018*	0.045***
	(0.008)	(0.017)	(0.010)	(0.013)
$D_{s,t} \times Post_t^{FR}$	0.006*	0.007*	0.004	0.005
	(0.003)	(0.004)	(0.003)	(0.004)
Observations	981,250	305,003	496,515	179,732
LSOA–Year fixed effects	Yes	Yes	Yes	Yes
Year–Quarter fixed effects	Yes	Yes	Yes	Yes
Property-level controls	Yes	Yes	Yes	Yes

Notes: This table examines whether the capitalization of SuDS differs before and after the introduction of the UK flood reinsurance scheme Flood Re introduced in 2016. The dummy $Post_t^{FR}$ equals one for transactions in calendar years $t \geq 2016$ and zero otherwise. Each column reports estimates from a two-way fixed-effects regression of the natural logarithm of transaction price on an indicator for SuDS presence within the LSOA, its interaction with the post-Flood Re period, LSOA fixed effects, calendar-year fixed effects, and the full set of property-level controls (type, vintage, tenure). Column (1) uses the full London sample; columns (2)–(4) restrict the sample to LSOAs in high-, moderate-, and low flash-flood risk bands, respectively, defined using Environment Agency AEP thresholds. Standard errors clustered at the LSOA level are reported in parentheses. Coefficients can be interpreted as semi-elasticities; for example, a coefficient β corresponds approximately to a $100 \times [\exp(\beta) - 1]$ percent change in price. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table E.6
Extensive- and intensive-margin effects of SuDS on housing prices.

Treatment variable	Extensive margin	Intensive margin		
		Continuous	Categorical	
	Binary SuDS indicator (1)	SuDS density (standardized) (2)	SuDS = 1 installation (3)	SuDS ≥ 2 installations (4)
Panel A. All LSOA				
SuDS treatment coefficient	0.038***	0.018***	0.029***	0.044*
	(0.008)	(0.006)	(0.010)	(0.025)
Observations	981,250	981,250	981,250	981,250
Panel B. High flash-flooding risk LSOA				
SuDS treatment coefficient	0.043***	0.022**	0.034**	0.055*
	(0.016)	(0.011)	(0.016)	(0.031)
Observations	305,003	305,003	305,003	305,003
Panel C. Moderate flash-flooding risk LSOA				
SuDS treatment coefficient	0.016*	0.011	0.015	0.024
	(0.009)	(0.007)	(0.010)	(0.020)
Observations	496,515	496,515	496,515	496,515
Panel D. Low flash-flooding risk LSOA				
SuDS treatment coefficient	0.046***	0.021**	0.031**	0.039*
	(0.012)	(0.010)	(0.015)	(0.022)
Observations	179,732	179,732	179,732	179,732
LSOA–Year fixed effects	Yes	Yes	Yes	Yes
Year–Quarter fixed effects	Yes	Yes	Yes	Yes
LSOA linear time trends	No	No	No	No
Population weights	No	No	No	No

Notes: The table reports two-way fixed-effects estimates of the extensive- and intensive-margin effects of SuDS on the natural logarithm of housing transaction prices. Column (1) reproduces results in Table 3 column (2) in which the treatment indicator equals one from the first calendar year in which any SuDS installation is completed within an LSOA and remains one thereafter. Column (2) replaces the binary indicator with a continuous intensity measure defined as the cumulative number of SuDS installations per hectare in each LSOA, standardized to have mean zero and unit standard deviation among treated LSOAs; coefficients in column (2) therefore represent the semi-elasticity of prices with respect to a one–standard-deviation increase in SuDS density. Columns (3) and (4) report estimates from a specification with mutually exclusive categorical treatment indicators for LSOAs that ever receive exactly one installation (SuDS = 1) and those that eventually host at least two installations (SuDS ≥ 2), with never-treated LSOAs as the omitted category. All models include LSOA and calendar-year fixed effects. Specifications are estimated for the full sample (Panel A) and separately by baseline flash-flooding risk status (Panels B–D), defined using Environment Agency AEP bands. Standard errors clustered at the LSOA level are reported in parentheses. Coefficients may be converted to percentage changes as $100 \times [\exp(\beta) - 1]$. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table E.7
Heterogeneity in SuDS capitalization by exposure in 2021 flood events.

	Subsample DiD estimates		Pooled interaction model
	Flood-exposed LSOA (1)	Non-exposed LSOA (2)	All LSOA (3)
SuDS treatment $D_{s,t}$	0.047***	0.036***	0.033***
	(0.014)	(0.010)	(0.009)
$D_{s,t} \times Post_t^{2021}$			0.006
			(0.007)
$D_{s,t} \times F_s \times Post_t^{2021}$			0.018**
			(0.008)
Observations	112,151	336,476	448,627
LSOA-year fixed effects	Yes	Yes	Yes
Year–Quarter fixed effects	Yes	Yes	Yes
Property-level controls	Yes	Yes	Yes
Sample window	2018–2024	2018–2024	2018–2024

Notes: This table examines whether capitalization of SuDS is stronger in neighborhoods that experienced realized flooding during the July 2021 flash-flood events. Flood exposure F_s is constructed from Sentinel–1 SAR VH backscatter over a 10 m buffer around the road network. For each LSOA, I compare mean backscatter in dry pre-event periods to the minimum value observed on 12 or 25 July 2021 and classify LSOAs in the top quartile of the resulting drop as "flood-exposed" ($F_s = 1$); all others are coded as "non-exposed" ($F_s = 0$). Columns (1) and (2) report separate two-way fixed-effects estimates of the average SuDS treatment effect for flooded and non-flooded LSOAs. Column (3) reports coefficients from a pooled interaction specification in which $D_{s,t}$ captures the average SuDS capitalization effect in non-flooded LSOAs before 2021, $D_{s,t} \times Post_t^{2021}$ captures the change in this effect after 2021 in non-flooded LSOAs, and $D_{s,t} \times F_s \times Post_t^{2021}$ captures the additional post-2021 capitalization in flooded LSOAs. All models include LSOA and calendar-year fixed effects, as well as property-level controls (type, vintage, tenure). The sample is restricted from 2018 to 2024. Standard errors clustered at the LSOA level are reported in parentheses. Coefficients may be converted to percentage changes as $100 \times [\exp(\beta) - 1]$. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

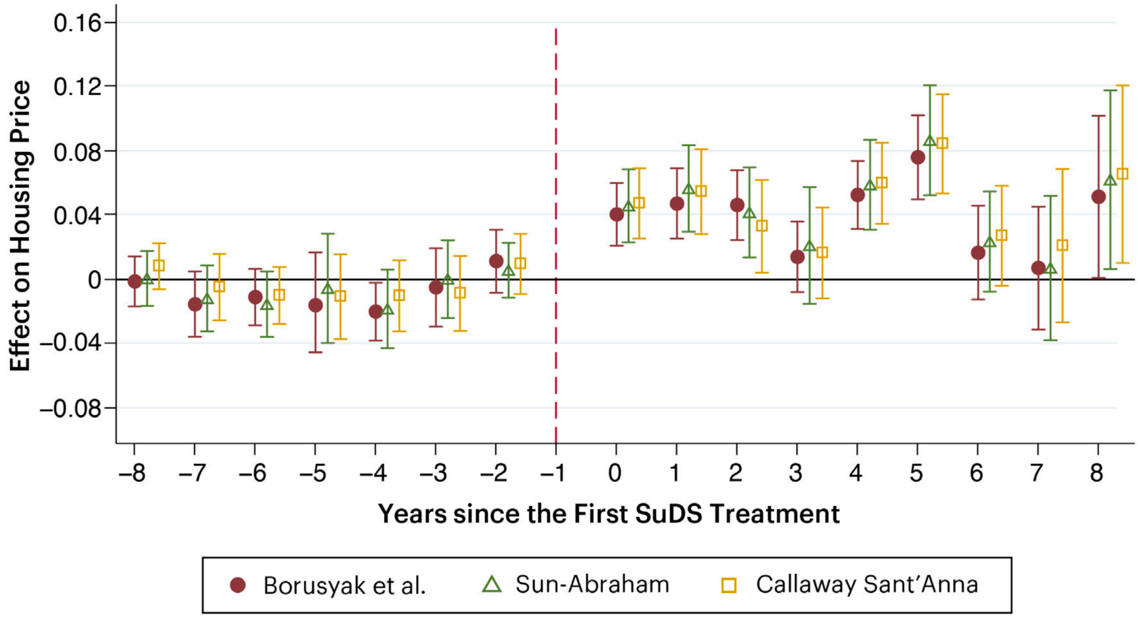

Fig. E.1. Event-study estimates of SuDS treatment effects under alternative estimators.
Notes: The figure plots dynamic treatment effects of SuDS implementation on the natural logarithm of housing transaction prices using three modern estimators for staggered adoption. The horizontal axis reports event time in years relative to the first SuDS delivery within an LSOA (with $t = -1$ omitted), and the vertical axis shows the estimated effect on the natural logarithm of prices. Circles report estimates from the imputation estimator of Borusyak et al. (2024); triangles report interaction-weighted estimates following Sun and Abraham (2021); squares report group–time average treatment effects following Callaway and Sant'Anna (2021). All specifications include LSOA and calendar-time fixed effects and the full set of property-level controls, and are estimated on the balanced 2010–2024 panel. Vertical bars denote 95 percent confidence intervals based on standard errors clustered at the LSOA level. The three estimators yield nearly identical pre-treatment paths and very similar post-treatment profiles, indicating that our main findings are not sensitive to the choice of staggered DiD estimator.

References

Abadie, A., 2005. Semiparametric difference-in-differences estimators. Rev. Econ. Stud. 72 (1), 1–19.

Angrist, J.D., Pischke, J.-S., 2009. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press.

Anguelovski, I., Connolly, J.J., Pearsall, H., Shokry, G., Checker, M., Maantay, J., Gould, K., Lewis, T., Maroko, A., Roberts, J.T., 2019. Why green "climate gentrification" threatens poor and vulnerable populations. Proc. Natl. Acad. Sci. 116 (52), 26139–26143.

Babovic, F., Mijic, A., 2019. The development of adaptation pathways for the long-term planning of urban drainage systems. J. Flood Risk Manag. 12 (S2), e12538.

Bae, J., Sohn, W., Newman, G., Gu, D., Woodruff, S., Van Zandt, S., Ndubisi, F., Wilkins, C., Lee, J., Tran, T., 2021. A longitudinal analysis of green infrastructure conditions in coastal texan cities. Urban For. Urban Green. 65, 127315.

Balaian, S.K., Sanders, B.F., Abdolhosseini Qomi, M.J., 2024. How urban form impacts flooding. Nat. Commun. 15 (1), 6911.

Beltrán, A., Maddison, D., Elliott, R., 2019. The impact of flooding on property prices: A repeat-sales approach. J. Environ. Econ. Manag. 95, 62–86.

Black, S.E., Denning, J.T., Dettling, L.J., Goodman, S., Turner, L.J., 2023. Taking it to the limit: Effects of increased student loan availability on attainment, earnings, and financial well-being. Am. Econ. Rev. 113 (12), 3357–3400.

Borusyak, K., Jaravel, X., Spiess, J., 2024. Revisiting event-study designs: robust and efficient estimation. Rev. Econ. Stud. 91 (6), 3253–3285.

Callaway, B., Sant'Anna, P.H., 2021. Difference-in-differences with multiple time periods. J. Econometrics 225 (2), 200–230.

Campos, L.C., Darch, G., 2015. Adaptation of UK wastewater infrastructure to climate change. Infrastruct. Asset Manag. 2 (3), 97–106.

Card, D., Krueger, A.B., 2000. Minimum wages and employment: a case study of the fast-food industry in new jersey and pennsylvania: reply. Am. Econ. Rev. 90 (5), 1397–1420.

Carpenter, C., Dobkin, C., 2011. The minimum legal drinking age and public health. J. Econ. Perspect. 25 (2), 133–156.

Chapman, C., Hall, J.W., 2022. Designing green infrastructure and sustainable drainage systems in urban development to achieve multiple ecosystem benefits. Sustain. Cities Soc. 85, 104078.

Chui, T.F.M., Liu, X., Zhan, W., 2016. Assessing cost-effectiveness of specific LID practice designs in response to large storm events. J. Hydrol. 533, 353–364.

Clement, M.A., Kilsby, C.G., Moore, P., 2018. Multi-temporal synthetic aperture radar flood mapping using change detection. J. Flood Risk Manag. 11 (2), 152–168.

Cohen, J.P., Field, R., Tafuri, A.N., Ports, M.A., 2012. Cost comparison of conventional gray combined sewer overflow control infrastructure versus a green/gray combination. J. Irrig. Drain. Eng. 138 (6), 534–540.

Cohen, J.P., Huang, Y., McMillen, D., 2025. Can fairness flow in philly? Real estate assessment equity for green stormwater infrastructure finance. J. Real Estate Financ. Econ. 1–25.

Contat, J., Hopkins, C., Mejia, L., Suandi, M., 2024. When climate meets real estate: A survey of the literature. Real Estate Econ. 52 (3), 618–659.

Costa, R., Bolte, E., Sharp, C., Bowers, C., 2024. Repeated and localized flooding is an underestimated challenge for urban disaster risk management. Nat. Cities 1 (9), 587–596.

Cousins, J.J., Hill, D.T., 2021. Green infrastructure, stormwater, and the financialization of municipal environmental governance. J. Environ. Policy Plan. 23 (5), 581–598.

Crompton, J.L., 2005. The impact of parks on property values: empirical evidence from the past two decades in the United States. Manag. Leis. 10 (4), 203–218.

Davis, J., Hagen, B., Mahid, Y., Pijawka, D., 2023. Evaluating property value impact from water-related 'green infrastructure': A hedonic modeling approach. J. Green Build. 18 (1), 3–16.

De Chaisemartin, C., d'Haultfoeuille, X., 2023. Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey. Econom. J. 26 (3), C1–C30.

Deryugina, T., Molitor, D., 2020. Does when you die depend on where you live? Evidence from hurricane katrina. Am. Econ. Rev. 110 (11), 3602–3633. http://dx.doi.org/10.1257/aer.20181026, URL https://www.aeaweb.org/articles?id=10.1257/aer.20181026.

Feng, Y., Nassauer, J.I., 2022. Aesthetics and the performance of stormwater infrastructure. In: Routledge Handbook of Urban Landscape Research. Routledge, pp. 163–180.

Finewood, M.H., Matsler, A.M., Zivkovich, J., 2019. Green infrastructure and the hidden politics of urban stormwater governance in a postindustrial city. Ann. Am. Assoc. Geogr. 109 (3), 909–925.

Fowdar, H.S., Neo, T.H., Ong, S.L., Hu, J., McCarthy, D.T., 2022. Performance analysis of a stormwater green infrastructure model for flow and water quality predictions. J. Environ. Manag. 316, 115259.

Galvin, E.M., BenDor, T.K., 2023. The economic impacts of green stormwater infrastructure: An evaluation of novel stormwater management policies in Washington, DC. Land Use Policy 134, 106915.

GLA, 2015. London Sustainable Drainage Action Plan: Draft for Public Consultation. Consultation Document, Greater London Authority, London, UK, URL https://www.london.gov.uk, Published by Greater London Authority, City Hall, The Queen's Walk, More London, London SE1 2AA.

Goodman-Bacon, A., 2021. Difference-in-differences with variation in treatment timing. J. Econometrics 225 (2), 254–277.
Haer, T., Botzen, W.W., de Moel, H., Aerts, J.C., 2017. Integrating household risk mitigation behavior in flood risk analysis: an agent-based model approach. Risk Anal. 37 (10), 1977–1992.
Hallegatte, S., Green, C., Nicholls, R.J., Corfee-Morlot, J., 2013. Future flood losses in major coastal cities. Nat. Clim. Chang. 3 (9), 802–806.
Herath, P., Bai, X., 2024. Benefits and co-benefits of urban green infrastructure for sustainable cities: six current and emerging themes. Sustain. Sci. 19 (3), 1039–1063.
Imai, K., Keele, L., Yamamoto, T., 2010. Identification, inference and sensitivity analysis for causal mediation effects. Statist. Sci. 25 (1), 51–71.
Jayasooriya, V., Ng, A., 2014. Tools for modeling of stormwater management and economics of green infrastructure practices: A review. Water, Air, Soil Pollut. 225 (8), 2055.
Johnson, D., Geisendorf, S., 2019. Are neighborhood-level SUDS worth it? An assessment of the economic value of sustainable urban drainage system scenarios using cost-benefit analyses. Ecol. Econom. 158, 194–205.
Jongman, B., 2018. Effective adaptation to rising flood risk. Nat. Commun. 9 (1), 1986.
Kaiser, Z.A., Akter, F., 2025. From risk to resilience and sustainability: Addressing urban flash floods and waterlogging. Risk Sci. 100011.
Kaźmierczak, A., Cavan, G., 2011. Surface water flooding risk to urban communities: Analysis of vulnerability, hazard and exposure. Landsc. Urban Plan. 103 (2), 185–197.
Kunreuther, H., Demidov, A., Pauly, M., Turcic, M., Wilson, M., 2023. Externalities in the wildland–urban interface: Private decisions, collective action, and results from wildfire simulation models for california. Risk Anal. 43 (5), 886–895.
Lamond, J., Proverbs, D., Hammond, F., 2010. The impact of flooding on the price of residential property: A transactional analysis of the UK market. Hous. Stud. 25 (3), 335–356.
Liu, W., Chen, W., Feng, Q., Peng, C., Kang, P., 2016. Cost-benefit analysis of green infrastructures on community stormwater reduction and utilization: A case of Beijing, China. Environ. Manag. 58 (6), 1015–1026.
Liu, L., Hudgens, M.G., Becker-Dreps, S., 2016. On inverse probability-weighted estimators in the presence of interference. Biometrika 103 (4), 829–842.
McLeod, M.K., Mickovski, S.B., 2024. The use of end-of-line SUDS for residential development. J. Sustain. Res. 6 (2), e240031.
Ossa-Moreno, J., Smith, K.M., Mijic, A., 2017. Economic analysis of wider benefits to facilitate SuDS uptake in London, UK. Sustain. Cities Soc. 28, 411–419.
Park, B., Kweon, B.-S., 2025. A systematic literature review of the economic effects of stormwater best management practices (BMPs) on housing prices. Landsc. J. 44 (1), 1–19.
Picon, A., 2014. Nature, infrastructure and cities. In: The Return of Nature. Routledge, pp. 172–180.
Pons, V., Abdalla, E.M.H., Tscheikner-Gratl, F., Alfredsen, K., Sivertsen, E., Bertrand-Krajewski, J.-L., Muthanna, T.M., 2023. Practice makes the model: A critical review of stormwater green infrastructure modelling practice. Water Res. 236, 119958.
Pryce, G., Chen, Y., Galster, G., 2011. The impact of floods on house prices: an imperfect information approach with myopia and amnesia. Hous. Stud. 26 (02), 259–279.
Rajapaksa, D., Zhu, M., Lee, B., Hoang, V.-N., Wilson, C., Managi, S., 2017. The impact of flood dynamics on property values. Land Use Policy 69, 317–325.
Rosen, S., 1974. Hedonic prices and implicit markets: product differentiation in pure competition. J. Political Econ. 82 (1), 34–55.
Skouralis, A., Lux, N., Andrew, M., 2024. Does flood risk affect property prices? Evidence from a property-level flood score. J. Hous. Econ. 66, 102027.
Sohn, W., Kim, H.W., Kim, J.-H., Li, M.-H., 2020. The capitalized amenity of green infrastructure in single-family housing values: An application of the spatial hedonic pricing method. Urban For. Urban Green. 49, 126643.
Sun, L., Abraham, S., 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. J. Econometrics 225 (2), 175–199.
Sunstein, C.R., 2018. The Cost-Benefit Revolution. MIT Press.
Trojanek, R., 2016. The impact of green areas on dwelling prices–the case of Poznań city. Entrep. Bus. Econ. Rev. 4 (2), 27–35.
Versini, P.-A., Kotelnikova, N., Poulhes, A., Tchiguirinskaia, I., Schertzer, D., Leurent, F., 2018. A distributed modelling approach to assess the use of blue and green infrastructures to fulfil stormwater management requirements. Landsc. Urban Plan. 173, 60–63.
Wu, Y., Yang, J., Li, Z., 2026. From vulnerability to adaptation: A systematic review of land-based economic outcomes of adverse climate events. Landsc. Urban Plan. 266, 105530.
Zenios, S.A., 2024. The climate-sovereign debt doom loop: what does the literature suggest? Curr. Opin. Environ. Sustain. 67, 101414.
Zuniga-Teran, A.A., Gerlak, A.K., Elder, A.D., Tam, A., 2021. The unjust distribution of urban green infrastructure is just the tip of the iceberg: A systematic review of place-based studies. Environ. Sci. Policy 126, 234–245.