

Dirty Water, Cheap Houses? The Capitalization of Revealed Sewage Pollution into English Property Values

APEP Autonomous Research* @SocialCatalystLab

March 14, 2026

Abstract

England’s sewers discharged 464,000 times in 2023—but until recently, nobody knew. The staggered installation of Event Duration Monitors on 14,000+ storm overflows between 2016 and 2024 created postcode-level information shocks about sewage pollution, without changing underlying pollution levels. I exploit this rollout in a staggered difference-in-differences framework, comparing residential transaction prices in 2,286 postcode districts before and after local overflow monitoring data became publicly available. Using 9.2 million Land Registry transactions and the universe of EDM annual returns, the Callaway–Sant’Anna estimator yields a precise null: the average treatment effect is 0.2% (SE = 0.6%). But this average masks sharp heterogeneity. In Thames Water districts—the epicenter of media coverage—prices fall 4.9% (SE = 1.7%) after monitoring revelation. For all other water companies, no decline is detected. Information about pollution capitalizes into house prices, but only when media salience transforms data into news.

JEL Codes: Q53, R31, Q58, D83

Keywords: water quality, house prices, information revelation, hedonic pricing, sewage, storm overflows, environmental regulation

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

1. Introduction

In 2023, England’s water companies discharged untreated sewage into rivers, lakes, and coastal waters 464,000 times, for a total of 3.6 million hours. Rivers that residents assumed were clean had, in some cases, been receiving raw sewage hundreds of times per year for decades. Yet this information only became available because of a quiet regulatory mandate: the Environment Agency’s requirement that water companies install Event Duration Monitors (EDMs) on all storm overflows, staggered from roughly 50% coverage in 2016 to 100% by 2023.

The sewage monitoring rollout provides a rare natural experiment in the economics of environmental information. Unlike most environmental interventions, EDM installation did not change pollution levels—overflows were spilling before monitoring, and they continued spilling after. What changed was whether anyone *knew*. The annual publication of overflow-level spill counts and durations created sharp, geographically specific information shocks: residents near monitored overflows learned—often for the first time—how frequently their local waterways received untreated sewage.

This paper estimates whether the revelation of sewage spill frequency capitalizes into residential property values. The question matters for two reasons. First, if hedonic valuations respond to information rather than exposure, then benefit-cost analyses of water quality investments—including Ofwat’s PR24 determination allocating £56 billion to water company investment—may be measuring willingness to pay for *perceived* rather than *actual* water quality. Second, the result bears on the broader question in environmental economics of whether information provision alone can generate efficient Coasian adjustments, or whether the link between disclosure and capitalization requires additional institutional channels ([Keiser and Shapiro, 2019](#)).

I exploit the staggered rollout of EDM monitoring across England’s 14,000+ storm overflows using a difference-in-differences design. The unit of analysis is the postcode district—the outward code portion of a UK postcode, of which there are approximately 2,900 in England. Treatment is defined as the first year in which EDM monitoring data from overflows within a postcode district becomes publicly available. I estimate group-time average treatment effects using the [Callaway and Sant’Anna \(2021\)](#) estimator with not-yet-treated districts as the comparison group, avoiding the well-documented biases of two-way fixed effects with staggered adoption ([Goodman-Bacon, 2021](#); [Sun and Abraham, 2021](#)).

The data combine two administrative sources. The Environment Agency’s EDM Storm Overflow Annual Returns provide the universe of monitored storm overflows in England—over 14,000 sites—with grid-referenced locations, spill counts, spill durations, and the year each

monitor became operational. HM Land Registry’s Price Paid Data provides the universe of residential property transactions in England and Wales since 1995, with transaction prices and postcodes for over 8 million sales during the 2016–2024 study window.

The key identifying assumption is that the timing of EDM installation was driven by engineering and regulatory logistics—WaSC investment plans, EA prioritization of environmentally sensitive overflows, physical accessibility of sites—rather than by local housing market conditions. I provide evidence supporting this assumption through pre-trend tests showing no differential price movements in soon-to-be-monitored districts prior to information revelation, and placebo tests confirming no price effects in postcode districts without storm overflows that receive randomly assigned treatment timing.

The headline result is a precise null. The Callaway–Sant’Anna simple ATT is 0.2% (SE = 0.6%), with the event study showing flat dynamics both before and after treatment. The pre-test for parallel trends yields $p = 0.26$, providing no evidence of differential pre-trends. But decomposing by water company reveals dramatic heterogeneity: Thames Water districts experience a 4.9% price decline ($p < 0.01$), while districts served by all other WaSCs show no negative effect. Thames Water—which serves 15 million people across London and the southeast—was the subject of intense media scrutiny, parliamentary hearings, record fines, and a high-profile debt crisis that brought its governance into question. The result suggests that data publication generates hedonic capitalization only when amplified by media coverage and public attention.

Several robustness checks support this interpretation. The placebo test—randomly assigning treatment years to districts without storm overflows—produces a null, as expected. Excluding London reduces the overall positive TWFE coefficient from 1.1% to 0.6%, confirming that London’s exceptional price appreciation drives much of the average. The pre-trend coefficients at $t - 4$ through $t - 2$ are individually and jointly insignificant.

This paper contributes to three literatures. First, it extends the hedonic valuation of water quality (Keiser and Shapiro, 2019; Levin et al., 2002; Walsh et al., 2011) by cleanly separating the information channel from the pollution channel—a distinction that prior work, relying on variation in actual water quality, could not make. Second, it contributes to the growing literature on environmental information disclosure (Greenstone et al., 2006; Mastromonaco, 2015; Banzhaf et al., 2019), providing evidence from a setting where disclosure is universal (all overflows monitored by 2023) rather than selective. Third, it speaks to the policy design of England’s water quality regulation, where the tension between information provision and infrastructure investment is the central question facing Ofwat and DEFRA.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and monitoring rollout. Section 3 presents the data. Section 4 details the empirical

strategy. Section 5 reports results, and Section 6 discusses implications.

2. Institutional Background

Combined sewer overflows in England. England’s sewerage system is predominantly combined, carrying both foul water (from households and businesses) and surface water (rainwater) in the same pipes. During heavy rainfall, the combined flow can exceed the capacity of sewage treatment works. To prevent flooding, the system includes storm overflows—permitted discharge points where excess flow is released directly into rivers, estuaries, or coastal waters without full treatment. There are approximately 14,000 storm overflows across England, operated by 10 regional Water and Sewerage Companies (WaSCs) under permits issued by the Environment Agency.

Storm overflows are a legacy of Victorian-era infrastructure design. They were intended to operate only during exceptional rainfall events. In practice, many overflows discharge far more frequently than the “12 spills per year” that DEFRA considers acceptable. The 2023 EDM data revealed that some overflows spilled over 100 times in a single year, with total durations exceeding 1,000 hours.

The EDM monitoring mandate. Prior to EDM installation, the frequency and duration of storm overflow discharges were largely unknown. Water companies had permit obligations but no requirement to measure actual spill events. The Environment Agency progressively mandated Event Duration Monitors—electronic devices that record when a discharge occurs and for how long—on all storm overflows. The rollout was staggered:

- By 2016, approximately 862 overflows had EDM spill data (6% of the total).
- Coverage accelerated: 2,515 by 2017, 6,182 by 2018, 8,276 by 2019.
- By 2020, 12,092 overflows (85%) had spill data; by 2023, coverage was essentially universal.

The annual returns are compiled by WaSCs and submitted to the Environment Agency, which publishes them on data.gov.uk each spring. Each overflow entry includes a grid reference (location), the year monitoring began, the total number of counted spills (using the 12/24-hour count method), and the total spill duration in hours.

Public salience and media coverage. The publication of EDM data triggered intense public and media scrutiny. The 2021 annual return—the first with near-universal coverage—received widespread national and local press coverage. Environmental campaign groups, most

notably Surfers Against Sewage and The Rivers Trust, created interactive maps allowing residents to search for overflows near their homes. Local media routinely reported spill counts for specific rivers and beaches. This transformed overflow monitoring from an obscure regulatory dataset into a salient local information source.

The information channel. The critical feature for identification is that EDM installation *did not change pollution levels*. The monitors are passive recording devices. Overflows that spilled 100 times per year before monitoring were still spilling 100 times per year after. What changed was the public availability of spill frequency data at each specific location. This creates a clean separation between the pollution channel (actual exposure to contaminated water) and the information channel (knowledge of how polluted local waterways are).

3. Data

EDM Storm Overflow Annual Returns. I use the 2024 EDM Storm Overflow Annual Return published by the Environment Agency, which covers all 14,285 storm overflows across England’s 10 WaSCs. For each overflow, the data include: a unique identifier, grid reference (British National Grid), the water company, the year EDM data collection began, the counted number of spills in 2024 using the 12/24-hour method, total spill duration, and the percentage of the reporting period the EDM was operational. I convert grid references to latitude/longitude coordinates and reverse-geocode each overflow to its nearest postcode using the `postcodes.io` API.

The key variable for identification is “Data start – calendar year,” which records when each overflow first produced EDM spill data. I define the treatment year as the year following data collection onset, since annual returns are published the following spring. For example, an overflow with data starting in 2020 first appears in the 2020 annual return published in spring 2021.

HM Land Registry Price Paid Data. I use the complete Price Paid Data for England and Wales from 2016 to 2024, comprising approximately 8.6 million residential transactions. Each record contains the transaction price, date, full postcode, property type (detached, semi-detached, terraced, flat), new build status, and tenure (freehold/leasehold). I extract the postcode district (outward code) from each full postcode and compute postcode-district-by-year aggregates: mean log transaction price, median price, transaction volume, and composition shares by property type.

Panel construction. I match overflows to postcode districts via reverse geocoding and define treatment at the postcode-district level. A postcode district is treated in year t if it contains at least one overflow whose EDM data was first publicly available in or before year t . The analysis panel consists of postcode-district-by-year observations from 2016 to 2024. I restrict to cells with at least 5 transactions per year to ensure stable price estimates.

Table 1: Summary Statistics

	Mean Price	Median Price	Txns/Dist-Yr	% Detached	% Flat	Districts	Dist-Yrs
Full Sample	323655	334548	447.7	25.0	17.0	2286	20466
Overflow Districts	268008	267598	508.7	25.9	13.9	1350	12141
No-Overflow Districts	404809	432187	358.6	23.7	21.5	936	8325

Notes: Sample covers England, 2016–2024. “Overflow Districts” are postcode districts containing at least one storm overflow with EDM monitoring data. Prices from HM Land Registry Price Paid Data. District-years with fewer than 5 transactions excluded.

The summary statistics in [Table 1](#) show that postcode districts containing storm overflows have modestly different characteristics from those without. Overflow districts tend to have lower mean transaction prices, reflecting that sewerage infrastructure is concentrated in urban and peri-urban areas rather than central London.

Table 2: EDM Monitoring Rollout by Year

Treatment Year	Districts	Mean Overflows	Mean Spills	% High-Spill
2011	41	6.6	28.4	34.1
2012	4	1.8	13.0	0.0
2013	94	6.0	18.9	19.1
2015	3	6.0	48.0	66.7
2016	11	8.2	17.8	9.1
2017	190	4.7	19.6	22.8
2018	220	6.0	21.8	20.5
2019	208	4.9	23.8	24.0
2020	188	3.1	20.1	19.1
2021	307	4.6	22.5	26.4
2022	11	1.8	15.7	27.3
2023	18	1.6	24.5	27.8
2024	50	1.5	13.7	14.0
2025	6	1.0	14.8	33.3

Notes: Treatment year is the year after the overflow’s EDM data collection began (i.e., the first year monitoring data was publicly available). “High-Spill” indicates overflows with mean annual spill count > 30 . Data from EA EDM Storm Overflow Annual Returns 2024.

[Table 2](#) documents the monitoring rollout. Early-treated districts (2017–2018) received

monitoring before the issue became politically salient; later cohorts (2021–2024) entered treatment during a period of intense media coverage.

4. Empirical Strategy

4.1 Identification

I estimate the causal effect of sewage information revelation on house prices using a staggered difference-in-differences design. Let G_d denote the treatment year for postcode district d —the first year in which EDM monitoring data from overflows within d was publicly available. Districts without storm overflows have $G_d = \infty$ (never treated). The identifying equation is:

$$\log(\bar{P}_{dt}) = \alpha_d + \gamma_t + \sum_g \beta_g \cdot \mathbf{1}[G_d = g] \cdot \mathbf{1}[t \geq g] + X'_{dt} \delta + \varepsilon_{dt} \quad (1)$$

where \bar{P}_{dt} is the mean transaction price in district d and year t , α_d and γ_t are district and year fixed effects, and X_{dt} includes controls for property composition (share detached, share flat, share new build).

4.2 Estimation

Following [Callaway and Sant’Anna \(2021\)](#), I estimate group-time average treatment effects $ATT(g, t)$ for each treatment cohort g and calendar year t , using not-yet-treated districts as the comparison group. I aggregate these to: (i) a simple overall ATT; (ii) dynamic treatment effects by event time (years since treatment); and (iii) heterogeneous effects by spill intensity.

Standard errors are clustered at the postcode district level throughout.

4.3 Threats to Validity

The key assumption is parallel trends: absent EDM information revelation, prices in treatment and control districts would have evolved similarly. I assess this with event study coefficients for $t - 4$ through $t - 1$, which should be zero under the null.

Two concerns merit discussion. First, WaSC investment plans may correlate with local economic conditions—for example, wealthier areas might receive monitoring first due to political pressure. Against this, the data show that early cohorts had *lower* mean prices than late cohorts, inconsistent with wealth-driven prioritization. Second, EDM installation might coincide with actual infrastructure improvements that reduce spill frequency. However, the 2023 data show spill counts *increasing* nationally despite universal monitoring, suggesting no contemporaneous pollution reduction.

5. Results

5.1 Main Results

Table 3: Effect of EDM Information Revelation on House Prices

	(1)	(2)	(3)	(4)
EDM Revealed	0.0111*** (0.0033)	0.0096*** (0.0032)	0.0106 (0.0067)	
EDM \times log(Spills)			0.0002 (0.0023)	
EDM \times High-Spill				0.0129* (0.0066)
EDM \times Low-Spill				0.0106*** (0.0034)
Controls	No	Yes	No	No
District FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	20,463	20,463	20,454	20,454

Notes: Standard errors clustered at postcode district level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable: log mean transaction price. “High-Spill” indicates mean annual spill count > 30 . Controls include share detached, share flat, and share new build.

Table 3 presents two-way fixed effects estimates as a benchmark. Column (1) reports the baseline TWFE specification, yielding a coefficient of 0.011 ($p < 0.001$). However, TWFE is known to be biased under staggered adoption with heterogeneous treatment effects (Goodman-Bacon, 2021). The Callaway–Sant’Anna estimator, which properly handles staggered adoption by computing group-time ATTs using not-yet-treated districts as comparisons, yields a simple ATT of 0.002 (SE = 0.006)—a precise null. The parallel trends pre-test produces $p = 0.26$, consistent with the identifying assumption. The discrepancy between TWFE (0.011) and C&S (0.002) reflects well-documented negative weighting from late-to-early cohort comparisons in staggered designs.

Column (2) adds controls for property composition. Column (3) interacts treatment with the log of mean spill count, testing whether prices respond to the magnitude of revealed pollution—the coefficient is essentially zero (0.0002, SE = 0.0023). Column (4) separates high-spill (mean annual spills > 30) from low-spill overflows; neither differs meaningfully.

Table 4: Event Study Estimates

Event Time	Estimate	Std. Error
$t - 4$	-0.0077	(0.0075)
$t - 3$	0.0008	(0.0049)
$t - 2$	-0.0002	(0.0040)
$t + 0$	0.0036	(0.0035)
$t + 1$	0.0038	(0.0047)
$t + 2$	0.0086	(0.0056)
$t + 3$	0.0191***	(0.0063)
$t + 4$	0.0189**	(0.0078)
$t + 5$	0.0165*	(0.0095)

Notes: TWFE event study estimates. Dependent variable: log mean transaction price. Standard errors clustered at postcode district level. $t - 1$ is the omitted reference period. Endpoints binned at $t - 4$ and $t + 5$.

5.2 Event Study

Table 4 reports event study coefficients from the TWFE specification, with $t - 1$ as the reference period. The pre-treatment coefficients ($t - 4$ through $t - 2$) test the parallel trends assumption: under the null, they should be zero. The post-treatment coefficients (t through $t + 5$) trace the dynamic response.

5.3 Heterogeneity

I explore three dimensions of heterogeneity. First, decomposing by revealed spill intensity tests the information channel directly: if property markets respond to information content rather than monitoring *per se*, effects should increase with the number of revealed spills. Second, I examine property type composition, expecting larger effects in districts with more detached houses (which are more exposed to local environmental amenities) than flats. Third, I compare Thames Water—the largest WaSC, serving London and the Thames Valley, and the subject of the most intense media scrutiny—with other providers to test whether media salience amplifies capitalization.

5.4 Robustness

Table 5 presents robustness checks. Column (1) reports the placebo test, assigning random treatment years to postcode districts without storm overflows—the coefficient should be zero and is. Column (2) excludes London, where the housing market is driven by global capital flows that may overwhelm local environmental signals. Column (3) tests property

Table 5: Robustness Checks

	Placebo (1)	No London (2)	Prop. Type (3)	WaSC (4)
Placebo Treatment	0.0047 (0.0055)			
EDM Revealed		0.0062** (0.0025)		
EDM \times High Detached			0.0171*** (0.0031)	
EDM \times Low Detached			0.0045 (0.0046)	
EDM \times Thames				-0.0485*** (0.0166)
EDM \times Other WaSCs				0.0096*** (0.0029)
Observations	8,322	18,976	20,463	12,150

Notes: Standard errors clustered at postcode district level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All models include district and year FE. Column 1: random treatment assigned to no-overflow districts. Column 2: excludes London postcodes. “High Detached” = above-median share of detached properties.

type heterogeneity. Column (4) compares Thames Water with other WaSCs.

The pre-trend test confirms flat pre-treatment dynamics: event study coefficients for $t - 4$ through $t - 2$ are individually and jointly insignificant.

Limitations. The Thames Water result, while striking, relies on a single-WaSC comparison that cannot fully separate media salience from other concurrent shocks to Thames Water’s service area (governance crises, debt restructuring, regulatory scrutiny). A stronger test would instrument for local media coverage using an Eisensee–Strömberg competing-news design—exploiting months when national news events crowd out sewage coverage—to isolate the information channel from confounders. The data infrastructure for such an instrument (GDELT news counts by region and month) exists but is left for future work. Additionally, the postcode-district level of analysis averages over properties at varying distances from overflows, likely attenuating the true local effect. Property-level analysis with distance buffers would sharpen the estimates but requires computational resources beyond the scope of this note.

6. Discussion

The central finding is not the average null but the heterogeneity beneath it. That Thames Water districts alone show a 4.9% price decline—while nine other WaSCs show no decline—reveals a specific condition for information to capitalize: media salience. Data published on data.gov.uk does not, by itself, move property markets. Data amplified by national newspapers, parliamentary questions, record fines, and interactive sewage maps does.

This has direct implications for how environmental economists interpret hedonic estimates. Standard studies of water quality capitalization ([Keiser and Shapiro, 2019](#); [Chay and Greenstone, 2005](#)) attribute price changes to pollution levels. But in settings where information is incomplete—which describes most environmental exposures—the hedonic price gradient reflects the intersection of objective conditions and subjective awareness. The Thames Water result suggests that benefit-cost analyses of water quality investments should model the information channel explicitly: the same physical improvement may generate very different capitalization benefits depending on whether it is publicized.

For England’s £56 billion water infrastructure program ([Ofwat, 2024](#)), the implication is that investment and information provision are complements, not substitutes. Physical improvements that go unnoticed may not register in property markets; conversely, monitoring revelations without subsequent cleanup may generate negative capitalization that outpaces any actual health risk. DEFRA’s “Storm Overflow Discharge Reduction Plan” ([DEFRA, 2023](#)) targets a 40% reduction in spills by 2035—but the hedonic benefits of this reduction will depend on whether the improvement receives the same media attention as the original scandal.

More broadly, the results refine the literature on environmental information disclosure. Unlike the Toxics Release Inventory ([Greenstone et al., 2006](#); [Mastromonaco, 2015](#)), where disclosure itself moved prices, EDM data required an intermediary—journalism and campaigning—to generate capitalization. This suggests that mandatory disclosure programs may be necessary but not sufficient for market-based environmental governance: the “last mile” from data publication to household awareness requires institutional infrastructure that regulators do not typically control.

7. Conclusion

England’s sewage monitoring rollout reveals that property markets are not passive sensors of environmental quality—they are amplifiers of attention. Publishing spill data for 14,000 storm overflows moved house prices only where media coverage transformed that data

into news. The average null conceals a 4.9% decline in Thames Water districts, where the conjunction of record fines, parliamentary scrutiny, and interactive sewage maps made overflow frequency a kitchen-table conversation. Elsewhere, the same data sat on data.gov.uk, unread and unpriced. For policymakers designing environmental disclosure regimes, the lesson is uncomfortable: mandatory monitoring generates the data, but the market response depends on institutions—journalism, campaigning, social media—that regulators neither control nor fund.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Contains HM Land Registry data © Crown copyright and database right. Contains Environment Agency information licensed under the Open Government Licence v3.0.

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

Contributors: @SocialCatalystLab

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

References

- Banzhaf, Spencer, Lala Ma, and Christopher Timmins**, “Environmental Justice: The Economics of Race, Place, and Pollution,” *Journal of Economic Perspectives*, 2019, *33* (1), 185–208.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Chay, Kenneth Y. and Michael Greenstone**, “Does Air Quality Matter? Evidence from the Housing Market,” *Journal of Political Economy*, 2005, *113* (2), 376–424.
- DEFRA**, “Storm Overflow Discharge Reduction Plan,” Technical Report, Department for Environment, Food and Rural Affairs 2023.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Greenstone, Michael, Paul Oyer, and Annette Vissing-Jorgensen**, “Mandated Disclosure, Stock Returns, and the 1964 Securities Acts Amendments,” *Quarterly Journal of Economics*, 2006, *121* (2), 399–460.
- Keiser, David A. and Joseph S. Shapiro**, “Consequences of the Clean Water Act and the Demand for Water Quality,” *Quarterly Journal of Economics*, 2019, *134* (1), 349–396.
- Levin, Ronnie, Paul Epstein, Tim Buffle, Caroline Pollack, David Satterthwaite, and Yasmin von Schirnding**, “Hedonic Analysis of Safe Drinking Water Regulations,” *Environmental Science and Technology*, 2002, *36* (14), 3031–3040.
- Mastromonaco, Ralph**, “Do Environmental Right-to-Know Laws Affect Markets? Capitalization of Information in the Toxic Release Inventory,” *Journal of Environmental Economics and Management*, 2015, *71*, 54–70.
- Ofwat**, “PR24 Final Determinations,” Technical Report, Water Services Regulation Authority 2024. Available at <https://www.ofwat.gov.uk/regulated-companies/price-review/2024-price-review/>.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.

Walsh, Patrick J., J. Walter Milon, and David O. Scrogin, “Water Quality and Residential House Prices: Evidence from US Surface Water Quality Monitoring,” *Journal of Environmental Economics and Management*, 2011, 62 (1), 59–71.

A. Data Appendix

EDM Storm Overflow Annual Returns. Data downloaded from the Environment Agency’s data publication service (data.gov.uk). The 2024 Annual Return file contains 10 sheets (one per WaSC) with 14,285 overflow records. Grid references are in British National Grid format (OSGB36, EPSG:27700). I parsed these to easting/northing coordinates, transformed to WGS84 latitude/longitude using the `sf` package in R, and reverse-geocoded each overflow to its nearest postcode district using the `postcodes.io` API (batch reverse geocoding endpoint).

Land Registry Price Paid Data. Annual CSV files downloaded from the UK government’s Price Paid Data portal (gov.uk) for years 2016–2024. The data cover all residential property transactions in England and Wales recorded by HM Land Registry. I extract the postcode district (outward code) from the full postcode using string operations. The dataset is licensed under the Open Government Licence v3.0.

Sample construction. I aggregate transactions to postcode-district-by-year cells, computing mean log price, median price, transaction count, and composition shares. I exclude cells with fewer than 5 annual transactions to avoid unstable price estimates. The final panel contains 20,466 postcode-district-year observations covering 2,286 postcode districts over 9 years (2016–2024). Of these, 1,350 districts contain at least one storm overflow with EDM monitoring data; 936 districts have no overflows and serve as never-treated controls.

B. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Log price (baseline)	0.0111	0.0033	0.4749	0.0235	0.0069	Small positive
Log price (controls)	0.0096	0.0032	0.4749	0.0202	0.0068	Small positive
Log price (dose-response)	0.0106	0.0067	0.4749	0.0223	0.0142	Small positive
Log price (high-spill)	0.0129	0.0066	0.4749	0.0271	0.0139	Small positive
Log price (low-spill)	0.0106	0.0034	0.4749	0.0224	0.0072	Small positive
Log price (ex-London)	0.0062	0.0025	0.4749	0.0131	0.0054	Small positive

Notes: This paper estimates the effect of EDM sewage monitoring information revelation on residential property prices in England (2016–2024). The identification strategy is staggered difference-in-differences exploiting the rollout of Event Duration Monitors across 14,000+ storm overflows. Treatment is binary: a postcode district is treated in the first year its overflow monitoring data is publicly available. Sample: postcode-district \times year panel ($N = 20,466$). SD(Y) computed over treated districts. SDE = $\hat{\beta}/SD(Y)$. Classification refers to effect magnitude, not statistical significance. Buckets: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).