Online Appendix for *Drinking Water Contamination and Home Prices: Evidence from California*

This Online Appendix provides supporting materials, additional technical details, and robustness checks for "Drinking Water Contamination and Home Prices: Evidence from California." Appendix Section A1 draws on Rosen (1974) to provide a theoretical framework for the empirical results presented in this study. Appendix Section A2 documents coverage of nitrate contamination issues in California from media and policymakers, establishing the salience of this issue for local residents and homebuyers. Appendix Section A3 provides a brief description of the Borusyak et al. (2024) imputation estimator used in the main analysis. Appendix Section A4 provides exact coefficients from the event-study results presented in Figure 3 from the main text and details the construction of CWS-level home price measures from ZIP-code level data. Appendix Section A5 demonstrates robustness of the main results to alternative estimators. Finally, Appendix Section A6 presents leave-one-out estimation results confirming that the findings from the main text are not driven by outlier water systems.

A1 Theoretical Framework

Following Rosen (1974), we can model home prices as reflecting the value of a set of bundled characteristics:

$$P = f(S, N, E) \tag{A.1}$$

Where P represents the transactions price of a given home, S denotes structural characteristics (square footage, bedrooms, etc.), N represents neighborhood attributes (schools, crime rates), and E captures environmental amenities including drinking water quality.

In the hedonic equilibrium, the marginal price of an attribute such as clean drinking water equals the marginal willingness to pay (*MWTP*) of the marginal buyer for that attribute. We can denote the marginal willingness to pay for water quality WQ as *MWTP(WQ)* and write:

$$\frac{\partial P}{\partial WQ} = MWTP(WQ) \tag{A.2}$$

In this setting, public notifications (PN) cause consumers to update their beliefs about perceived water quality from WQ_0 to WQ_1, with $WQ_1 < WQ_0$, for homes served by the affected CWS. Assuming consumers were unaware of contamination issues prior to notification (or underestimated the severity of those issues by assuming that they had been resolved), the information provided by the PN causes consumers to revise their assessment of the amenity value of housing located in the affected CWS. At the new equilibrium post-PN, the equilibrium price for housing will be lower, reflecting both reduced buyer willingness to pay and the response of sellers to this demand shift. This represents a local partial-equilibrium around the prevailing hedonic price schedule.

The static difference-in-differences (DiD) coefficient reported in the main results represents the average change in home prices associated with a notification:

$$\mathbb{E}[\Delta P | Notification] \approx f(S, N, WQ_1) - f(S, N, WQ_0) \tag{A.3}$$

Year fixed effects absorb common market shocks, and CWS fixed effects absorb time-invariant unobserved amenities and baseline differences across systems. Under the parallel trends assumption, this represents the causal effect of new information about water contamination. From the perspective of the Rosen model, this effect may represent the capitalized value of clean drinking water, subject to the assumption that the composition of transacted homes does not shift, and that PNs do not induce supply shocks or spatial spillovers. Importantly, because *MWTP* varies across consumers and PNs may induce sorting across locations, the estimated price effects represent equilibrium responses that combine changes in both the composition of buyers and their willingness to pay, rather than necessarily identifying the *MWTP* of any particular consumer type.

Finally, this overall framework relies on the assumption that Tier 1 nitrate PNs produce salient shifts in perceived water quality, which is reasonable given the acute health risks such contamination poses, and the attention that nitrate contamination has received from media and policymakers as described in Appendix Section A2 below.

A2 Nitrate Contamination in California

> *The fight for clean water in California is a fight for justice. No Californian should be exposed to toxins or hazardous waste because of where they live, and no child should get sick because of where he or she goes to school. Now is the time to invest in drinking and wastewater systems that can keep our communities healthy and safe.*
>
> *– CA State Senator Bill Monning and Tom Steyer (Monning and Steyer, 2017)*

Nitrate contamination from agricultural fertilizers and animal waste has long stood as one of California's most pressing water quality challenges, affecting millions of residents across the state. In 2008, the severity of nitrate contamination concerns in California prompted legislation requiring the State Water Board to conduct a wide-ranging investigation into the causes and costs of nitrate contamination, as well as potential remediation options. The ensuing investigation generated an 8-part series of reports, conducted jointly with the University of California (Groundwater Nitrate Project, 2012). The authors of this study noted that nitrate contamination was a widespread issue throughout California and that "nitrate problems will likely worsen for several decades."

A 2020 study by the Environmental Working Group (EWG) found that nitrate contamination issues had not abated in the intervening years following the Groundwater Nitrate Project. The EWG report found that elevated levels of nitrates have been reported in CWS across California and that contamination levels have risen over the past two decades for many of these systems (Schechinger, 2020). A report by Becker (2024) for *CalMatters* described nitrate contamination as one of the most pervasive water quality issues facing the state. This report noted that unsafe drinking water has left residents of rural areas affected by water contamination, "afraid to drink tap water, or even bathe their children in it, relying on bottled water instead." An article by Del Real (2019) in *The New York Times* described the extent of nitrate contamination issues in California's agricultural areas, and echoed the findings from Becker (2024), noting that many residents relied on purchasing bottled water. High-profile op-eds discussing state water quality issues published in *Los Angeles Times* (Leslie, 2017) and *The Mercury News* (Monning and Steyer, 2017) also specifically cite nitrate contamination as a key concern.

As noted in Appendix Section A1 above, for nitrate PNs to affect home prices, local residents and homebuyers must be cognizant of such notifications. In this context, the attention that nitrate contamination issues have generated among policymakers, the media, and local residents of affected areas underscores the salience of these issues for prospective homebuyers. The adverse health consequences of drinking contaminated water, and the costs and inconvenience of mitigation efforts such as purchasing bottled water, suggest that PNs regarding nitrates represent clear disamenities for prospective homebuyers, which may be capitalized into home prices. More generally, the widespread nature of nitrate contamination issues in California, and the disruptions they cause for affected communities, underscores the policy relevance of this study.

A3 Imputation Estimator

The results reported in the main text are derived using the imputation estimator developed by Borusyak et al. (2024), referred to below as the BJS estimator. This estimator addresses concerns about the estimation of event-study and difference-in-differences (DiD)-style research designs via Ordinary Least Squares (OLS). In this section, I provide a brief summary of important features of the BJS estimator.

BJS estimation is a two-stage process in which we start by using not-yet and never-treated observations to estimate fixed effects ϕ_w and γ_t and generate predicted values of counterfactual $Y_{wt}(0)$ for treated observations. We then construct treatment effects estimates as weighted average of differences between Y and $\hat{Y}$ for treated observations. As with any estimator in this setting, credible causal inference with the imputation-based approach requires an identifying assumption about the counterfactual trajectory of Y in the absence of treatment. Mapping the standard parallel trends assumption to the imputation setting, we are assuming that the predicted values $\hat{Y}_{wt}(0)$ reflect the trend treated CWS would have followed in absence of treatment. In the event-study results presented in Figure 3 in the main text, I report the p-value from a joint significance test of the pre-treatment event time coefficients, which allows us to assess the plausibility of this assumption.

A4 Construction of CWS-Level Home Price Measures

I gather data on community water systems (CWS) from California's State Water Resources Control Board. These records include GIS files with geographic boundaries for CWS service areas, in addition to CWS-level records on total service population (California State Water Resources Control Board, 2025). ZIP-code level ZHVI home price records are merged with CWS-level service boundaries using ZIP Code Tabulation Areas (ZCTA) provided by the Census Bureau. In conjunction with the ZCTA boundary files, I also gather Census data on ZCTA population counts. The primary outcome variable in my analysis is CWS-level weighted average home values. In order to make the ZIP-code level home values provided by the ZHVI representative of home prices across a given CWS (which may span multiple ZIP codes), I adopt an approach that allows me to account for both the geographic overlap of ZCTAs and CWS as well as the population distribution across ZCTAs.

I construct a weighted average home value for each CWS by first conducting a spatial merge of all ZCTAs and CWS. For each intersecting ZCTA j and CWS k pair, I calculate the proportion of the total CWS area that is covered by that matched ZCTA, denoted f_{jk}, given by the following:

$$f_{jk} = \frac{Area(ZCTA_j \cap CWS_k)}{Area(CWS_k)} \tag{A.4}$$

So that f_{jk} represents the geographic coverage fraction of CWS k by ZCTA j. To ensure proper weighting across multiple ZCTAs that may intersect with a single CWS, I normalize these coverage fractions:

$$w_{jk}^{geo} = \frac{f_{jk}}{\sum_{j'} f_{j'k}} \tag{A.5}$$

Where w_{jk}^{geo} is the normalized geographic weight for ZCTA j and CWS k, ensuring that weights sum to 1 for each CWS, accounting for the presence of other ZCTAs j' that intersect with CWS

k in addition to ZCTA j.

Using the ZCTA-level population data, I am able to construct weights based on both geographic intersection and population. Denoting the population of ZCTA j as Pop_j^{ZCTA}, I calculate a population-adjusted weight as:

$$w_{jk}^{pop} = \frac{f_{jk} \times Pop_j^{ZCTA}}{\sum_{j'}(f_{j'k} \times Pop_{j'}^{ZCTA})} \tag{A.6}$$

The final CWS-level home price measures are calculated as:

$$\overline{HP}_k^{geo} = \sum_j w_{jk}^{geo} \times HP_j \tag{A.7}$$

$$\overline{HP}_k^{pop} = \sum_j w_{jk}^{pop} \times HP_j \tag{A.8}$$

Where HP_j is the ZHVI home price for ZCTA j. Finally, as an additional check on the robustness of the results from the main text, I use ZCTA boundary definitions for both 2000 and 2010 because ZCTA boundary definitions may change over time.

In the results presented in Table A1 below, I show that the baseline event-study results presented in the main text are robust to variations in the specific approach taken to constructing weighted-average home prices. The first column of Table A1 reports estimates corresponding to the event-study and static DiD depicted results in Figure 3 of the main text. The estimates here use population-weighted average home values $\overline{HP}_k^{pop}$ and 2010-vintage ZCTA boundary definitions. In Column (2) I estimate the same event-study using $\overline{HP}_k^{geo}$ with 2010-vintage boundary definitions and no population weighting; finally, in Column (3) I show results using $\overline{HP}_k^{geo}$ with 2000-vintage ZCTA boundary definitions. Across all three specifications, the results are not meaningfully altered by the choice of ZCTA boundary vintage or weighting approach, with static difference-in-differences estimates that are within $\sim$0.5 percentage points of one another and event-study coefficients that likewise follow similar patterns in the pre- and post-periods.

A5 Robustness to Alternative Estimator Selection

In this section, I employ two alternative, modern DiD estimators in order to demonstrate the robustness of the results presented in the main text. I use both the OLS-based strategy proposed by Wooldridge (2021) as well as the two-stage estimation strategy proposed by Gardner et al. (2025). The Wooldridge estimator uses extended two-way fixed effects to allow for heterogeneous treatment effects across cohorts and time periods. Similar to the imputation estimator used in the main text, the two-stage approach by Gardner et al. first estimates fixed effects using only untreated/not-yet-treated observations, then computes treatment effects from residualized outcomes. Both of these estimators are designed to be robust to concerns about traditional OLS-based estimation for the research design employed in this study, where treatment timing is staggered and treatment effects may be heterogeneous.

I present the results from using these alternative specifications in Table A2 below. In the first column, I report the baseline estimate from Appendix Table A1 as a comparison point. In Column (2), I show the results from estimating the same specification using the Gardner two-stage approach, while in Column (3), I show the results obtained using the Wooldridge estimator. Across all three columns, the estimates are qualitatively very similar, indicating that the conclusions derived from the baseline specification employed in the main text are not sensitive to the choice of estimator.

A6 Leave-One-Out Estimation

In this section, I demonstrate that the results presented in the main text are not driven by outlier CWS by conducting a leave-one-out cross-validation exercise. This robustness check provides an assessment of whether any individual CWS exerts outsize influence on the baseline aggregate treatment effect estimate. I iterate over each CWS that is included in the analysis sample and which experiences at least one nitrate PN over the sample time period. In each iteration, I remove a given CWS from the analysis sample and re-estimate the static DiD coefficient reported in Appendix Table A1, Column (1). By sequentially excluding each treated unit from the estimation sample, this procedure allows me to examine the stability of the main findings across different sample compositions.

After iterating over all CWS that have experienced a nitrate PN, I am able to construct and plot the distribution of estimated static DiD coefficients and their corresponding t-statistics. These distributions provide visual evidence of the degree of sensitivity of the main findings to sample composition. If a small number of outlier CWS with larger-than-average responses to nitrate PNs are driving the results, then we would expect this distribution to exhibit dispersion or skewness in iterations where those influential CWS are excluded. Such a pattern would raise concerns about the generalizability of the findings. However, if most CWS tend to experience similar home price declines following PNs, suggesting a systematic market response to such notifications, then we would expect a more concentrated distribution centered around the full-sample estimate.

The results from this robustness exercise are plotted in Figure A1. I show the distribution of estimated static DiD coefficients in Panel A, and the distribution of t-statistics in Panel B. In the first panel, the distribution of treatment effect estimates is highly concentrated, with point estimates clustered around the coefficient of -0.0576 reported in Appendix Table A1, Column (1), with an average value of -0.058. This stability across sample compositions provides evidence that the documented relationship between nitrate PNs and home prices represents a systematic pattern rather than an artifact of outlier CWS. Likewise, the distribution of t-statistics is centered around -2.68, consistently remaining well beyond the conventional two-sided 95% confidence threshold plotted in gray at 1.96. This consistency of statistical significance across all iterations further demonstrates that the documented effects are not dependent on the inclusion of any particular CWS in the estimation sample, thereby strengthening confidence in both the internal validity and external validity of the main findings.

References

Becker, R. (2024). I won't let them drink the water: The California towns where clean drinking water is out of reach. *CalMatters*. September 10. Available online.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *The Review of Economic Studies*, 91(6):3253–3285.

California State Water Resources Control Board (2025). California drinking water system area boundaries. GIS data, Division of Drinking Water. Available online.

Del Real, J. A. (2019). They grow the nation's food, but they can't drink the water. *New York Times*. May 21. Available online.

Gardner, J., Thakral, N., To, L. T., and Yap, L. (2025). Two-stage differences in differences. Working paper, University of Mississippi, Department of Economics. Available online.

Groundwater Nitrate Project (2012). Addressing nitrate in California's drinking water. Technical report, Center for Watershed Sciences UC Davis.

Leslie, J. (2017). California's water crisis is dangerous, just like Flint's. Will the state clean it up once and for all? *Los Angeles Times*. May 4. Available online.

Monning, B. and Steyer, T. (2017). Opinion: Tax water to provide clean water to Californians whose supply is contaminated. *The Mercury News*. August 14. Available online.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55.

Schechinger, A. W. (2020). EWG investigation: Across farm country, nitrate pollution of drinking water for more than 20 million Americans is getting worse. Technical report, Environmental Working Group.

Wooldridge, J. M. (2021). Two-way fixed effects, the two-way Mundlak regression, and difference-in-differences estimators. Working paper, Michigan State University, Department of Economics. SSRN 3906345.

Table A1: *Baseline Event-Study Results from Figure 3 of Main Text and Robustness of Results to Alternative Mappings of CWS and ZCTAs*

	(1)	(2)	(3)
Panel A: Static DiD Estimates			
Nitrate Public Notification	-0.0576***	-0.0548**	-0.0605***
	(0.02)	(0.02)	(0.02)
N	2875	2875	2875
Panel B: Event-Study Estimates			
$t = -2$	-0.0066	-0.0091	-0.0146
	(0.01)	(0.01)	(0.02)
$t = -1$	-0.0041	-0.0069	-0.0134
	(0.02)	(0.02)	(0.02)
$t = 0$	-0.0154*	-0.0169*	-0.0209**
	(0.01)	(0.01)	(0.01)
$t = 1$	-0.0392***	-0.0401***	-0.0462***
	(0.01)	(0.01)	(0.01)
$t = 2$	-0.0504***	-0.0506***	-0.0566***
	(0.02)	(0.02)	(0.02)
$t = 3$	-0.0401**	-0.0396**	-0.0447**
	(0.02)	(0.02)	(0.02)
$t = 4$	-0.0356	-0.0342	-0.0399*
	(0.02)	(0.02)	(0.02)
N	2409	2409	2409
ZCTA Boundary Definition Year	2010	2010	2000
Population Weights	Yes	No	No
P-Value from Pre-Trends Test	0.83	0.80	0.64

Notes: Data is aggregated to the CWS-year level for the years 2000 to 2024. All results are generated using the robust BJS imputation estimator described in Appendix Section A3 with standard errors clustered at the CWS level. The *p*-value of a joint significance test of the pre-treatment coefficients is reported in the bottom row of the table. The estimates from Column (1) correspond to the baseline event-study output reported in Figure 3 of the main text. In Column (1), weighted average home prices at the CWS level are calculated using 2010-vintage ZCTA boundaries and weighting by ZCTA-level population; in Column (2), weighted averages are constructed using 2010-vintage ZCTA boundaries and no population weights; in Column (3), weighted average home prices are calculated using 2000-vintage ZCTA boundaries and no population weights.

***p<0.01, **p<0.05, *p<0.10

Table A2: *Robustness of Results to Alternative Estimators*

	BJS	Gardner	Wooldridge
Treatment Effect	-0.0576***	-0.0688***	-0.0576**
	(0.0215)	(0.0244)	(0.0263)
N	2875	2875	2875

Notes: Data is aggregated to the CWS-year level for the years 2000 to 2024. Column (1) reports the baseline static DiD estimate reported in Appendix Table A1, Column (1) above, derived using the BJS estimator described in Appendix Section A3. In Columns (2) and (3), I generate comparable static DiD estimates based on the Gardner et al. estimator and the Wooldridge estimator, respectively, both of which are briefly described in Appendix Section A5. All specifications include CWS and year fixed effects with standard errors clustered at the CWS level.
***p<0.01, **p<0.05, *p<0.10

Figure A1: *Leave-One-Out Estimation Results*

Panel A: *Distribution of Leave-One-Out Estimated Treatment Effects*

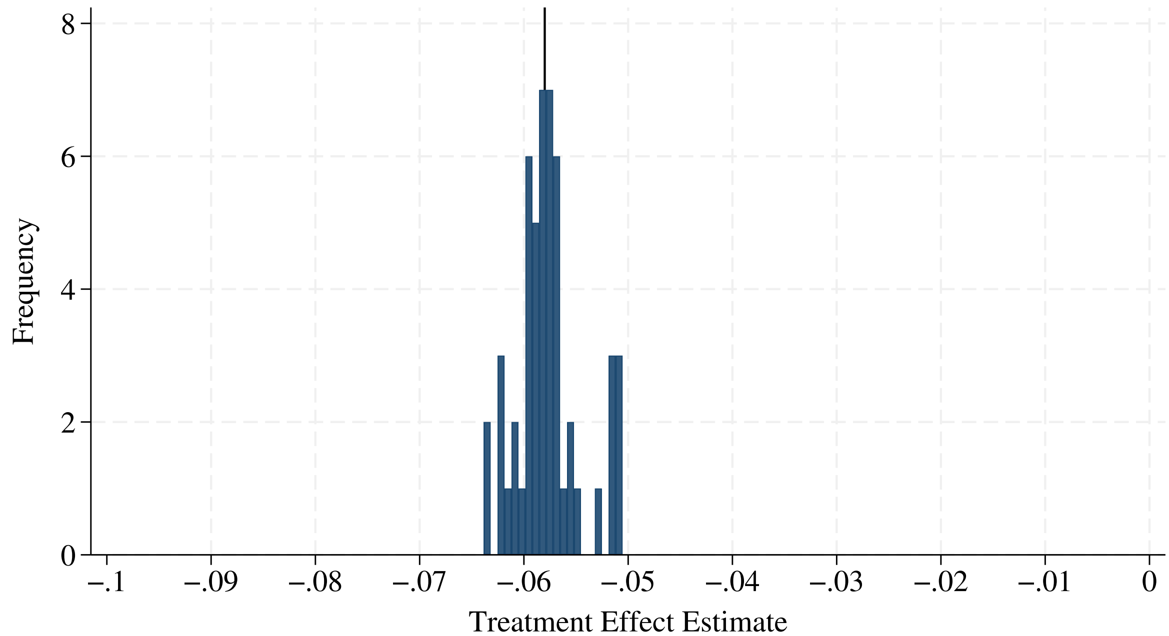

Median coefficient estimate of −0.058 denoted with black line.

Panel B: *Distribution of Leave-One-Out t-Statistics*

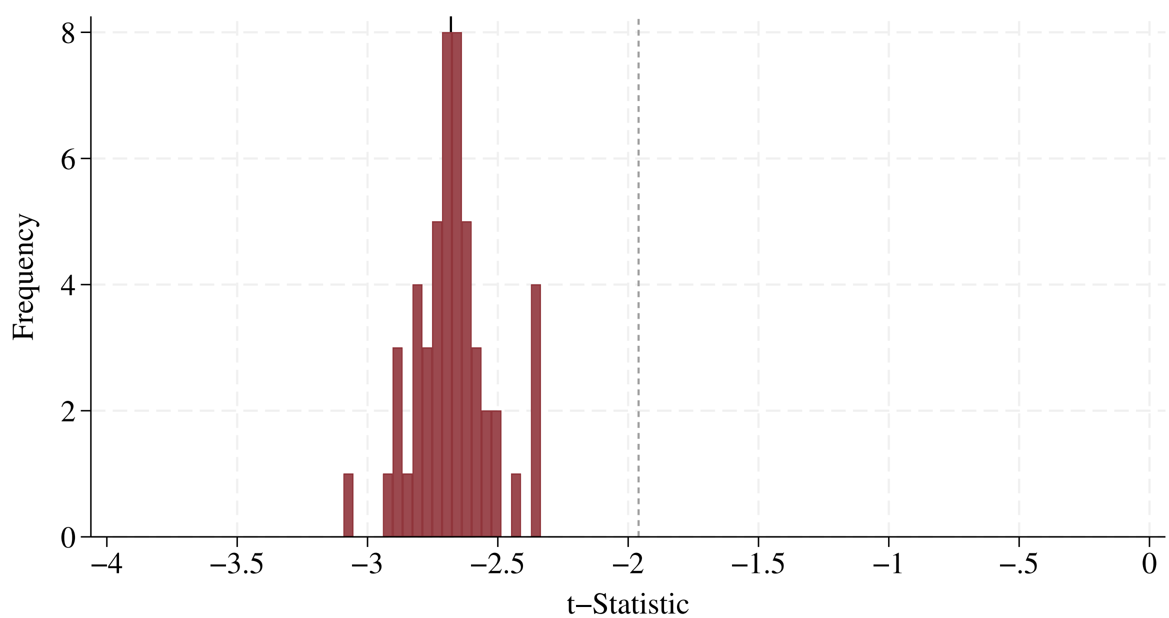

Median t−statistic of −2.68 denoted with black line. Gray line at $x = -1.96$ denotes two−sided 5% significance level.