

ONLINE APPENDIX

Inefficiencies in Local Infrastructure: Evidence from Drinking Water in California

Pierre Bodéré, Giulia Brancaccio, and Karam Kang

August 19, 2026

CONTENTS

OA1. Data Construction	2
OA1.1. Construction of Investment Measures	2
OA1.2. Drinking Water Quality Data	4
OA1.3. Constructing a System-Specific Surface-Water Pollution Index	6
OA1.4. Construction of Local Election Variables	9
OA2. Effect of Investment on Water Quality	11
OA3. Estimation of Policymakers' Parameters	11
OA3.1. Estimation details	11
OA3.2. Model fit	13
OA4. Counterfactual Results and Implementation Details	14
OA4.1. Distributional incidence of accountability	14
OA4.2. Counterfactual water rate increases without subsidies	15
References	20

OA1. DATA CONSTRUCTION

This appendix provides additional details on the construction of the main data used in the analysis. We first describe how we combine administrative financing records to measure drinking water infrastructure investment. We then discuss the construction of our drinking water quality measures and the system-specific surface-water pollution index used in the empirical analysis. Finally, we describe how we construct measures of local election contestation and election timing.

OA1.1. Construction of Investment Measures. This section describes how we construct measures of drinking water infrastructure investment by combining multiple administrative data sources and classifying the purpose of each financing instrument.

OA1.1.1. Data Sources and Investment Aggregation. We collect information on government debt, loans, and subsidies used to finance water infrastructure investments from four sources: (i) outstanding debt records from California’s Financial Transactions Reports (FTR), (ii) debt issuance records from the California Debt and Investment Advisory Commission (CDIAC), (iii) capital contributions from federal and state governments reported in water enterprise income statements in the FTR, and (iv) Drinking Water State Revolving Fund (DWSRF) loans and grants obtained through FOIA requests.

For each outstanding debt record in the FTR data, we observe the amount, issuance year, and a short description of the project purpose. We use this information to determine whether the debt corresponds to new financing or refinancing and whether it is related to drinking water infrastructure.¹ We combine these records with CDIAC data to improve coverage and measurement, as CDIAC provides more detailed information on debt characteristics and issuance.

After identifying all relevant sources of financing—including bonds, federal and state loans, DWSRF loans and grants, and direct subsidies—we compute the total amount of new funding raised for water infrastructure investment in each jurisdiction-year. Because investment financing may occur over multiple consecutive years, we group contiguous years with positive financing into an “investment spell.” We assign the first year of each spell as the investment year and sum total financing over the spell as the investment amount. Finally, we define an investment indicator equal to one if the total investment amount exceeds \$500,000 (in 2020 dollars) in that spell.

¹We exclude refinancing activities from our investment measure.

OA1.1.2. *Classification of Investment Purpose.* A key challenge in constructing investment measures is distinguishing drinking water infrastructure from other types of public investment. We classify municipal debt records into three categories based on the stated purpose of the debt and, when available, the nature of the pledged revenue: drinking water infrastructure, wastewater or sewer infrastructure, and unrelated projects. Drinking water infrastructure includes projects related to water supply, treatment, distribution, storage, contamination mitigation, and capital improvements supporting potable water systems. Wastewater or sewer infrastructure includes projects related to wastewater, sewer, sludge, or effluent systems. All other projects—including public safety, transportation, solid waste, recreation, and general government expenditures—are classified as unrelated.

We assign each record to one of these categories using a keyword-based decision rule applied to project descriptions and revenue fields. Specifically, if the description or revenue field contains wastewater or sewer-related terms, the record is classified as wastewater infrastructure. If not, but the description contains keywords associated with unrelated projects, the record is classified as unrelated. If the description instead contains primary drinking water keywords—such as references to water supply, treatment technologies, storage, distribution, or contaminants—the record is classified as drinking water infrastructure. When the description contains only more general, secondary water-related terms (such as system improvements or construction), the record is classified as drinking water infrastructure only if the pledged revenue field indicates a water-related funding source; otherwise, it is classified as unrelated. Records that do not meet any of these criteria are classified as unrelated.²

The keyword classification relies on predefined sets of terms. The first set concerns wastewater/sewer:

WASTEWATER, WASTE WATER, SEWER, SEWERS, EFFLUENT, SLUDGE, WW

The second set is for exclusion:

POLICE, FIRE, LIGHT, STREET LIGHT, ANIMAL, DOG, ELECTRIC, POWER, VEHICLE, TRUCK, COPIER, PRINTER, MAIL, COMPUTER, SERVER, PENSION, AMBULANCE, DISPATCH, RADIO, CAR, FLEET, GARAGE, HOUSING, PARK, LIBRARY, AIRPORT, HARBOR, PORT, SOLID WASTE, TRASH, LANDFILL, CEMETERY, RECREATION, SPORT, PLAYGROUND, TENNIS, GOLF, MUSEUM,

²For special districts, which are primarily water-focused entities, we do not distinguish between primary and secondary keywords and classify all water-related keywords as drinking water infrastructure.

PHONE, POSTAGE, HONEYWELL, PUMPER, INFUSION PUMP, SOLIS PUMP, HOWELL, ENERGY, SEAWELL, FARM, CNG TANK, GAS TANK, BOAT, HUD, SECTION 108, MARINA, ECONOMIC DEVELOPMENT, LIFT STATION, MAIN BUILDING, DIGESTER, BIOSOLID, COMM DEV, COPY MACHINE, OFFICE

Next, the set of primary water keywords includes:

WATER, DESAL, DESALINATION, TANK, WELL, GROUND WATER, SAFE DRINK, DRINK, DAM, PIPE, CHLORIN, FILTRAT, FILTER, PURIFY, PURIFICATION, METER, PUMP, RESERVOIR, MEMBRANE, ULTRAFILTRAT, OSMOSIS, DISINFECT, BACKWASH, CLEARWELL, BRINE, EQUALIZATION TANK, ARSENIC, LEAD, PFAS, PFOA, PFOS, NITRATE, NITRITE, MANGANESE, IRON, CONTAMIN, HEXAVALENT, CHROMIUM, MCL, BUREAU OF RECLAMATION, CENTRAL VALLEY PROJECT, DELTA, HYDROELECTRIC, SRF, DWSRF, STATE REVOLVING FUND, SWRCB, SCADA, CATHODIC, CORROSION, POTABLE, AWT, BASIN, ARROYO

The secondary water keywords are:

SYSTEM, IMPROV, CAPITAL, FACILIT, CONSTRUCT, REMEDIAT, REMOV, ABAT, DISTRIBUT, INFRASTRUCT, INTAKE, BOOSTER, AERATION, TREAT, EXPAND, UPGRAD, REPLAC, EXPANS, CONNECT, PILOT, ACQUI, AUGMENT

Lastly, the keywords on water revenues are

WATER, NET WATER, WATER FEES, WATER REVENUE, WATER SYSTEM, WATER FUND, WATER CHARGES, WATER NET

In applying these rules, wastewater indicators take precedence over all other classifications, and exclusion keywords override any water-related keywords. Primary water keywords are sufficient to classify a project as drinking water infrastructure without requiring confirmation from the revenue field, whereas secondary keywords require such confirmation. Subsidy data, including DWSRF loans and grants and capital contributions reported in the FTR, are restricted to drinking water investments by construction and therefore do not require classification.

OA1.2. Drinking Water Quality Data.

OA1.2.1. California Water Quality Analyses Data. Our main source on drinking water pollution is the Water Quality Analyses data maintained by the California State

Water Resources Control Board. We combine recent records from the Safe Drinking Water Information System (SDWIS, post-2011) with historical records from the Water Quality Management system (WQM, 1974–2010). Under the 1974 Safe Drinking Water Act (SDWA), the federal government sets drinking water standards, while states are responsible for monitoring public water systems, inspecting compliance, and maintaining records (Allaire et al., 2018). In California, public water systems collect samples at multiple stations and submit them to laboratories certified under the Environmental Laboratory Accreditation Program (ELAP). These ELAP-certified laboratories report test results directly to the California Division of Drinking Water (DDW), and these reports constitute our primary raw data.

Each record contains the sampled public water system, sample collection date, laboratory identifier, analyte name and code, test result, and reported unit of measurement. We undertake several harmonization steps across reporting systems and over time. We begin by restricting attention to contaminants that can generate health-based violations under the National Primary Drinking Water Regulations (NPDWR). For most such contaminants, the regulation specifies a maximum contaminant limit (MCL), usually expressed in milligrams per liter, though radionuclides are reported in picocuries per liter or millirems per year. One notable exception is fecal coliform, for which the rule is based on repeat-sample procedures rather than an MCL. Because California does not maintain electronic records of coliform samples, we exclude this contaminant from the analysis. We also exclude contaminants regulated through treatment techniques rather than MCLs, such as viruses and turbidity. After these restrictions, our analysis covers 74 federally regulated pollutants spanning four groups: inorganic chemicals, radionuclides, organic chemicals, and disinfection byproducts. Finally, as investments in our analysis could impact any part of the drinking water treatment cycle in a system, from water source to treatment plant, we keep all samples from every type of sampling station.

Many sampling records are reported as non-detects, indicating that the concentration falls below the detection limit of the sampling method. Following Keiser et al. (2023), we set these observations to zero. We then convert all readings into the units in which the relevant MCL is defined. Finally, again following Keiser et al. (2023), we aggregate the data to the system–pollutant level by computing the share of readings within each year-jurisdiction that exceed the corresponding MCL. To account for seasonal variation in contamination and differences in sampling schedules across systems, we also construct, for each system–pollutant pair, the fraction of readings

TABLE OA1. Federally Regulated Contaminant Readings

Contaminant type	Number of readings	Number of pollutants	Share zero	% above MCL	Median readings pollutant-system-year
Inorganic Chemicals	1,880,703	16	0.538	4.42	3.0
Radionuclides	156,158	4	0.367	7.71	3.0
Disinfection Byproduct	1,569,156	4	0.631	0.35	16.0
Organic Chemicals	5,327,577	50	0.961	0.84	3.0

Notes: This table provides descriptive statistics of the contaminant readings data for public water systems matched to jurisdictions in our sample over the time period 2001-2023.

collected in each month. As reported in Table OA1 below, after this cleaning, the readings data comprises 74 pollutants. The median pollutant is sampled multiple times per entity, year. Broadly, the number of pollutants per chemical type, and the fraction of readings above maximum concentration limits are consistent with Keiser et al. (2023), except for disinfection byproducts where our data has a smaller share of readings exceeding the limit.

OA1.2.2. SDWA Violations. Our second source of data for drinking water pollutions are the drinking water violations from the SDWA “Violations and Enforcements” data set available online. This dataset records violations at the public water system level. Violations can be of different type, including Monitoring, Reporting, and MCL violations. We focus on MCL violations only, ignoring other health-based violations such as Maximum Residual Disinfectant Level and Treatment Technique violations. Violations can be of different severity which translates into different public notification tiers, as specified by the Public Notification Rule (2000, revised in 2006). This rule ensures that the public is informed of drinking water violations posing potential risks to public health. The most serious category, Tier I, requires immediate notice within 24 hours after the public water system learns of the violation. Tier I public notifications might be issued via TV, radio, hand delivery, or any other method to reach all persons served by the public water system. Tier II requires public notification “as soon as practical” within 30 days, typically via mail. In addition, drinking water systems are required to issue annual consumer confidence reports (CCR) as per the 1998 CCR rule.

OA1.3. Constructing a System-Specific Surface-Water Pollution Index. We leverage monitoring data from the Environmental Protection Agency’s STORET database to construct a measure of ambient surface-water pollution that is relevant for

the drinking-water quality of each public water system. The idea is that surface-water contamination directly hinders the delivery of clean drinking water: higher contaminant concentrations in untreated source water require more intensive treatment (Price and Heberling, 2018).

STORET stations are distributed across California’s streams, rivers, lakes, and reservoirs and report contaminant concentrations, sampling locations, and sampling dates. The central measurement challenge is determining which ambient readings are informative about pollution risk at a given drinking-water system. A simple assignment based on a fixed spatial radius or temporal window is unlikely to capture this relationship accurately. We do not observe the precise intake locations or source-water bodies used by each system, and contaminants move through hydrologic networks through flow, dispersion, dilution, and chemical or biological transformation. Consequently, the relevant monitoring stations and temporal lags are not known *ex ante*. The relationship may also differ by source-water type: systems relying on surface water may be affected more directly by nearby or upstream readings, whereas contamination affecting groundwater systems may emerge over longer distances or time horizons.

We therefore construct an aggregate pollution index by training an attention-based neural network (Fig. OA1). Specifically, our aggregate pollution index is built to predict PWS-by-month-by-pollutant MCL violation probabilities using nearby ambient surface-water readings. For each drinking-water observation, the network aggregates a set of STORET readings observed at different stations, chemicals, distances, and lags. The attention mechanism allows the model to learn which readings are most predictive, rather than imposing a fixed spatial radius or temporal window. The model learns source-type-specific distance and lag penalties, allowing surface-water and groundwater systems to place different weights on nearby ambient readings. The architecture also allows for cross-chemical learning: surface-water readings for one pollutant can predict drinking-water violation risk for another pollutant through learned STORET-chemical and target-chemical embeddings. We then use the cross-fitted predicted violations as our aggregate pollution index, which we then leverage as an instrument for realized drinking-water pollution. Because predictions are generated out of fold at the PWS level, the model used to construct the predictions does not rely on any data from the target system for training. Figure OA3 shows the first stage of the regression.





Figure OA1: Spatial-Temporal Surface to Drinking Water Pollution Model.

Notes: This figure describes the architecture of the drinking water pollution prediction model. **Drinking water readings** define the prediction unit b : a PWS, a target chemical, and a month, together with the source category. The associated outcome is the drinking-water violation indicator y_b , defined based on readings regulated exceeding maximum concentration limit. The main predictors are **surface water readings** are a set of nearby STORET tokens indexed by k , each with features $\mathbf{x}_{bk}$, chemical identity c_{bk}^S . Each surface water reading k , relative to a system reading b , is characterized by and time-space distances (d_{bk}, ℓ_{bk}) that will determine the weight placed on that reading in aggregation. This aggregation is performed by the attention layer. Learnable parameters are shown in bold, including $\mathbf{W}^{\text{tok}}, \mathbf{W}^a, \mathbf{W}^v, \mathbf{W}^{\text{out}}$, and the attention-decay parameters $\boldsymbol{\lambda}_b^d, \boldsymbol{\lambda}_b^\ell$.

FIGURE OA2. Training and Validation Losses



Notes: This figure shows the training and validation losses across epochs during the fit of the model represented in Figure OA1. Training is performed in 5 folds splitting the data by drinking water systems. To predict on the left-out fold representing 20% of systems, a model trained on the remaining 80% is used.


OA1.4. Construction of Local Election Variables. This section describes the construction of the local election variables used in the analysis. We begin by identifying elections for the governing officials responsible for water infrastructure decisions and use the occurrence of an observed election to measure whether a jurisdiction holds a contested election in a given two-year period. We then classify observed elections as on-cycle or off-cycle according to whether they coincide with regularly scheduled statewide elections.

OA1.4.1. Identifying (contested) elections. We identify elections for the governing officials responsible for water infrastructure decisions by combining records from the California Elections Data Archive with election information collected from county election offices. The relevant offices depend on the jurisdiction's governance structure. For cities, we use mayoral and city council elections. For independent special districts, which are governed by directly elected boards, we use elections for district board members. Dependent special districts do not elect their own governing boards; we therefore link each district to its parent county government using the Financial Transaction Report data and use elections for the county board of supervisors.

When defining election frequency, we aggregate the election records to two-year periods. This is because we expect at least one seat to come up for election in each two-year period.³ However, an election may not appear in the election records when no seat is contested. For example, when the number of candidates does not exceed the number of available seats and the candidates assume office without an election.

³Although an individual officeholder may serve a four-year term, these governing bodies contain multiple seats whose terms are generally staggered.

FIGURE OA3. First Stage by System Water Source



Notes: This figure shows binscatter plots illustrating the first stage of the instrument constructed from the predicted probabilities of violation by chemical-month produced by the model presented in Figure OA1. The pollution measure and instrument are residualized on drinking water system and year fixed effects. The left panel shows the overall correlation, the middle panel shows the correlation for systems with groundwater source, the right panel shows the correlation for systems with surface water source.

We classify a jurisdiction-period as having a contested election if we observe at least one election for a relevant governing office during that period. When no such election is observed, we code the jurisdiction-period as having no contested election.

OA1.4.2. *Election timing: On-cycle vs. off-cycle.* We classify each observed local election according to whether it coincides with a regularly scheduled California statewide election. Specifically, an election is coded as on-cycle when its date exactly matches the date of a presidential or gubernatorial general election or a regularly scheduled statewide primary election. We obtain statewide election dates from the California Secretary of State's historical election records.⁴ We do not classify statewide recall elections or other special elections as on-cycle elections, because these elections do not occur according to the regular statewide electoral calendar.

⁴<https://www.sos.ca.gov/elections/prior-elections/statewide-election-results>.

We aggregate the election-level records into two-year electoral cycles, pairing each odd-numbered year with the subsequent even-numbered year and labeling the cycle by the even year. For each local government and two-year cycle, we calculate the fraction of its observed elections that were held on-cycle. Thus, the resulting measure equals one when all observed elections in the cycle coincided with regular statewide elections, zero when none did, and a value between zero and one when the jurisdiction held both on- and off-cycle elections.

OA2. EFFECT OF INVESTMENT ON WATER QUALITY


Figure OA4 presents the dynamic effects of investment on water pollution, using specification (11), for different sub-samples of the data: below-25th percentile and above-75th percentile groups for a given criterion (investment size, system size, and climate) or three groups depending on the source of water. We do not find evidence for a large heterogeneity in the effects by the amount of the investment (Panel A), the size of the system as measured by the number of connections (Panel B), the source of water (groundwater or surface water; Panel C), or the climate as represented by the average annual amount of rain during the period of study (Panel D).

OA3. ESTIMATION OF POLICYMAKERS' PARAMETERS

This appendix provides additional details on the estimation of the policymaker-side parameters and the fit of the estimated model. We first describe how we solve the policymaker's dynamic investment problem and estimate its parameters by maximum likelihood. We then assess model fit by comparing predicted and observed investment patterns along two key dimensions: the cost of investment and expected drinking-water pollution.

OA3.1. Estimation details. Recall from Section 6.5 that we estimate $\tilde{\pi}(\cdot)$ and θ_{λ_2} from turnout and election data, respectively, outside the model. The remaining policymaker parameters are $\boldsymbol{\psi} \equiv (\theta_{\lambda_0}, \theta_{\lambda_1}, \theta_{\kappa})$. For a given parameter vector $\boldsymbol{\psi}$, we solve the policymaker's dynamic program by value function iteration on a discretized (l^r, l^q) endogenous state space representing rate and expected pollution investment stock variables. These stocks are repleted when investment occurs and depreciate over time, consistent with the shape of our event studies estimated in Figure 1. To reduce the computational complexity, we assume that policymakers have perfect foresight over the evolution of exogenous state variables, $\boldsymbol{x}_{jt}$, during our sample, and expect

FIGURE OA4. Investment Effectiveness Across Water System Types



Notes: Each plot presents the coefficient estimates of the dynamic effects of investment on water pollution, using specification (11), for two sub-samples of the data: below-25th percentile and above-75th percentile groups, based on the amount of the investment (Panel A), the size of the system as measured by the number of connections (Panel B), the source of water (ground water or surface water; Panel C) and the average annual amount of rain during the period of study (Panel D).

them to remain fixed at their 2022 values thereafter.⁵ In particular, the exogenous determinants of pollution and water rate, α_{jt} and γ_{jt} , respectively, are included in those exogenous state variables over which the policymaker has perfect foresight, and thus are not integrated over. In each endogenous state (l^r, l^q) , we first construct residents' utility using the estimated willingness-to-pay $u(q_{jt}, r_{jt}; \omega(y_i))$ and aggregate resident utilities into the policymaker's flow payoff associated with elections using the

⁵We have also considered an alternative assumption that policymakers do not anticipate future changes in $\mathbf{x}_{jt}$ and obtain similar results. Both approaches provide a parsimonious way to focus on the dynamics of rate and pollution induced by investment while still capturing local differences in demographics and amenities. Explicitly modeling beliefs over the evolution of the many variables contained in $\mathbf{x}_{jt}$ would substantially expand the state space.

income distribution among those that turn out:

$$\Pi_{jt}(l^r, l^q; \theta_{\lambda_0}, \theta_{\lambda_1}) = \sum_y \lambda(\mathbf{x}_{jt}; \theta_{\lambda_0}, \theta_{\lambda_1}, \hat{\theta}_{\lambda_2}) \hat{\pi}_{jt}(y) u(q_{jt}(l^q), r_{jt}(l^r); \hat{\omega}(y)),$$

where $q_{jt}(l^q)$ and $r_{jt}(l^r)$ denote the mappings from the grid states (l^q, l^r) into current pollution and rates. The law of motion for the discretized state variables implied by our event study estimates are described in Equations (3) and (5). Given these primitives and the private cost of investing $\kappa(\mathbf{x}_{jt}; \theta_\psi)$, we solve the optimal policy of each jurisdiction, using value function iteration in our terminal period, 2022, and then backward induction for earlier periods. We thus recover the choice-specific value difference $\Delta V_{jt}(l^r, l^q; \boldsymbol{\psi})$ and implied conditional choice probabilities

$$P(a_{jt} = 1 \mid l^r, l^q, \mathbf{x}_{jt}; \boldsymbol{\psi}) = \frac{\exp(\Delta V_{jt}(l^r, l^q; \boldsymbol{\psi}))}{1 + \exp(\Delta V_{jt}(l^r, l^q; \boldsymbol{\psi}))},$$


which we evaluate at observed states and use to construct the log-likelihood

$$\begin{aligned} \ell(\boldsymbol{\psi}) = & \sum_{j,t} a_{jt} \log P(a_{jt} = 1 \mid l_{jt}^r, l_{jt}^q, x_{jt}; \boldsymbol{\psi}) \\ & + (1 - a_{jt}) \log (1 - P(a_{jt} = 1 \mid l_{jt}^r, l_{jt}^q, x_{jt}; \boldsymbol{\psi})). \end{aligned}$$

We maximize the average log-likelihood using gradient-based optimization, with automatic differentiation and GPU acceleration using pytorch.

OA3.2. Model fit. Figure OA5 evaluates the model’s ability to reproduce key empirical investment patterns. Panel (A) examines the relationship between investment probability and the magnitude of water rate increases associated with investment. In the data, investment becomes less likely as the implied cost of investment rises, reflecting the higher financial burden on residents. The model closely matches this declining relationship, indicating that the estimated cost component of the policy-maker’s payoff captures the observed sensitivity of investment decisions to financing costs. Panel (B) examines how investment probability varies with expected pollution levels. Both the data and the model show that investment is more likely when expected pollution is higher, reflecting greater expected benefits from improving water quality. Overall, the figure shows that the estimated model successfully reproduces two central empirical patterns: investment is less likely when costs are high and more likely when environmental need is greater.

FIGURE OA5. Model Fit on Investment Patterns




Notes: This figure shows the model fit along two specific dimensions. Panel (A) shows that the model reproduces the decreasing relationship observed in the data between investment probability and the magnitude of water rate increases from investment. Panel (B) shows that the estimated model successfully reflects that investment is more likely to occur in system-years where expected pollution is likely to be higher, for exogenous reasons such as weather and the water source used.

OA4. COUNTERFACTUAL RESULTS AND IMPLEMENTATION DETAILS

This appendix provides additional results and implementation details for the counterfactual exercises presented in Section 8. We first examine the distributional incidence of greater accountability by showing how the resulting changes in investment vary with jurisdictions' demographic and socioeconomic characteristics. We then describe how we construct the counterfactual increases in household water rates under no financial support from the federal and state governments, combining the estimated effects of investment on water rates with information on project-level financing and municipal borrowing terms.

OA4.1. Distributional incidence of accountability. Figure OA6 relates the jurisdiction specific change in investment under the high-accountability counterfactual to several demographic characteristics. The investment reallocation induced by greater accountability is concentrated in socioeconomically disadvantaged communities. Specifically, investment increases are larger in jurisdictions with lower shares of college-educated and White residents, higher shares of Hispanic residents, lower homeownership rates, and lower median household incomes.

FIGURE OA6. Distributional Incidence of Greater Accountability



Notes: Each panel relates a jurisdiction's percentage change in average investment probability under the high-accountability counterfactual, relative to the baseline, to the indicated demographic characteristic. In the high-accountability counterfactual, each jurisdiction's accountability index is set to the level associated with consistently contested elections, while all other model primitives and the existing subsidy environment are held fixed. Demographic characteristics are transformed using $\log(1 + x)$. The solid lines show linear fits, and the reported $\hat{\beta}$ values are the corresponding slope coefficients.

These patterns do not arise because the counterfactual directly targets jurisdictions based on their demographic composition. Rather, all jurisdictions are assigned the same high-accountability benchmark, and the resulting investment response depends on their local costs and welfare gains from improving drinking-water quality. The figure therefore shows that the investment shortfalls generated by weak accountability are disproportionately concentrated in disadvantaged communities. Combined with the larger pollution reductions among high-pollution systems documented in Figure 4, these results imply that strengthening accountability would make the allocation of drinking-water infrastructure investment more progressive.

OA4.2. Counterfactual water rate increases without subsidies. Our estimated effect of investment on water rates reflects the financing arrangements observed in the data, including both direct subsidies and below-market government lending. To evaluate the role of these types of federal and state financial support, we construct the increase in household water rates that would be required to finance an infrastructure

investment in the absence of grants and subsidized loans. We first present the accounting framework and then describe how we estimate its components using project-level financing and municipal debt data.

OA4.2.1. *Framework.* This section develops a simple accounting framework to construct the water rate increase households would have experienced in the absence of federal or state support. We denote by K the investment amount, which, under the observed financing environment, may be covered by a combination of market-rate borrowing, subsidized government loans, and grants. In the no-subsidy counterfactual, we assume that the same financing requirement K is instead financed entirely at market terms. Subsidies therefore affect households' water rates by changing the annual debt-service burden associated with financing K . The framework's key result is that the unsubsidized rate increase for households equals the observed, subsidized rate increase scaled up by the ratio of debt-service payments without and with subsidies.

Suppose the infrastructure project requires external financing K at interest rate i^B for T^B periods with fully amortizing level payments. The implied annual debt service payment is

$$C^{\text{unsub}}(K, i^B, T^B) = Ka(i^B, T^B), \quad (1)$$

where the annualization factor $a(\cdot, \cdot)$ converts one dollar of principal into the level annual payment implied by the interest rate and maturity as follows

$$a(i, T) = \frac{i}{1 - (1 + i)^{-T}}. \quad (2)$$

Let N denote the number of residential ratepayers. If the current average annual household water bill is r , the proportional water rate increase required to finance the investment in the absence of government support, denoted α^{unsub} , is

$$\alpha^{\text{unsub}} = \frac{C^{\text{unsub}}(K, i^B, T^B)}{rN}. \quad (3)$$

Grants eliminate repayment; subsidized loans at lower interest rates and grant longer maturity. Let s^B denote the share of investment costs financed via bonds issued by the local government with interest rate i^B and maturity T^B , and let s^L be the share of investment costs financed through subsidized loans, at the interest rate of i^L with maturity T^L . The remaining share, $1 - s^B - s^L$, is covered by grants. The corresponding total cost of investment is

$$C^{\text{base}}(K, s^B, s^L, i^B, i^L, T^B, T^L) = K(s^L a(i^L, T^L) + s^B a(i^B, T^B)), \quad (4)$$

where grant funding does not enter, because it creates no repayment obligation for the local water system. The associated increase in water rates for households, α^{base} , can therefore be written as

$$\alpha^{\text{base}} = \frac{C^{\text{base}}(K, s^B, s^L, i^B, i^L, T^B, T^L)}{rN}, \quad (5)$$

Combining (3) and (5) yields the counterfactual rate increase that would occur in the absence of external financial support:

$$\alpha^{\text{unsub}} = \alpha^{\text{base}} \times \frac{C^{\text{unsub}}(K, i^B, T^B)}{C^{\text{base}}(K, s^B, s^L, i^B, i^L, T^B, T^L)}. \quad (6)$$

Intuitively, this adjustment rescales the observed rate increase by the ratio of debt-service payments with and without subsidies.⁶

OA4.2.2. Estimation. To compute α^{unsub} using Equation (6) requires two inputs: the baseline rate increase α^{base} and the ratio of annualized debt-service costs with and without subsidies. We recover the first, $\hat{\alpha}^{\text{base}}$, directly from our difference-in-differences estimates. For the second input, we start by computing $C^{\text{base}}(K, i^B, i^L, T^B, T^L, s^B, s^L)$ and $C^{\text{unsub}}(K, i^B, T^B)$ for observed investments. This requires estimating the interest rate and maturity of a market-rate revenue bond, as well as the terms of a subsidized government loan. We denote the resulting borrowing terms by $(\hat{i}^B, \hat{i}^L, \hat{T}^B, \hat{T}^L)$.⁷ Next, we estimate separate Poisson pseudo-maximum-likelihood models for $C^{\text{unsub}}(K, \hat{i}^B, \hat{T}^B)$ and $C^{\text{base}}(K, s^B, s^L, \hat{i}^B, \hat{i}^L, \hat{T}^B, \hat{T}^L)$. Each model includes the logarithm of households served, the logarithm of median household income, the logarithm of the prevailing household water rate, and year fixed effects.⁸ We use the fitted conditional means to predict the annualized cost of an investment in every jurisdiction-year. Below, we describe in detail how we recover borrowing terms (i^B, i^L, T^B, T^L) for realized investments.

Revenue bond terms. We use debt-level records from the California Debt and Investment Advisory Commission (CDIAC), which report each bond’s coupon rate,

⁶The argument above assumes that households bear the full cost of investment. In practice, treated water also serves commercial and industrial users, so households bear only a share of these costs. Equation (6) remains valid as long as the household share of investment costs is unchanged by the subsidy.

⁷The FTR data provide the size of each investment and its financing composition, while borrowing terms—interest rates and maturities—come from the separate CDIAC database. Because matching individual financing records across the two sources is not straightforward, we estimate these borrowing terms rather than matching them directly.

⁸Our recovered annualized costs of realized projects are highly right-skewed. We therefore winsorize each cost measure at the first and ninety-ninth percentiles, when estimating Poisson regressions.

maturity, source of repayment, and purpose. To reduce selection concerns, we estimate financing terms using all municipal bonds in the CDIAC data rather than restricting the sample to drinking-water debt. We regress the log of maturity on indicators for the source of repayment and project purpose, together with government, demographic, and water-system characteristics. We similarly regress the log of the coupon rate on these variables and year fixed effects. We then predict the maturity and coupon rate of a hypothetical revenue bond issued for a drinking water project in each jurisdiction where we observe an investment.⁹

Subsidized government loan terms. We construct the expected terms of subsidized government loans for every observed investment using records from the Drinking Water State Revolving Fund (DWSRF), which we obtained from our FOIA request. Although there are other subsidized government loan programs run by the state government, we assume that subsidized loans' terms are similar to those of the DWSRF program. Because repayment periods are concentrated at 20 years and the available sample of loans is relatively small, we set the maturity of a DWSRF loan to $\hat{T}_{jt}^L = 20$. DWSRF loans may carry either the program's posted interest rate or a zero interest rate. We estimate the probability of receiving a zero-interest loan as a function of indicators for whether the system falls below the program's income thresholds and whether it qualifies as a small system. We then construct the expected DWSRF coupon rate as


$$\hat{i}_{jt}^L = \left(1 - \widehat{\Pr}(\text{zero interest} | \mathbf{x}_{jt})\right) i_t^{\text{DWSRF}},$$

where i_t^{DWSRF} is the posted DWSRF interest rate in year t , which is available at California State Water Resources Control Board's website.

OA4.2.3. *Estimated Financing Costs and Subsidy Adjustment.* Figure OA7 summarizes our estimates of how government support lowers infrastructure financing costs and how these estimates are used to construct the no-subsidy counterfactual. Panel (A) compares the predicted annualization factors, as defined in (2), for municipal bonds and subsidized government loans. The distribution for subsidized loans is concentrated around 5–6 percent, whereas the distribution for municipal bonds is

⁹Since not all local governments in our analysis have bond records in the CDIAC database, we do not include government fixed effects. Although debt size predicts both maturity and the coupon rate, including the logarithm of principal changes the R^2 only modestly: from 0.706 to 0.739 in the maturity regression and from 0.713 to 0.714 in the coupon-rate regression. We therefore exclude principal from the specifications used to predict financing terms.

FIGURE OA7. Financing Costs with and without Government Support



Notes: Panel (A) plots the distributions of the predicted annualization factors for municipal bonds and subsidized government loans, where bond terms are predicted from CDIAC debt records and subsidized-loan terms are based on DWSRF program rates, the predicted probability of receiving a zero-interest loan, and an assumed 20-year maturity. Panel (B) plots the predicted ratio of the no-subsidy financing cost to the baseline financing cost. The vertical dashed line denotes the median. This ratio is used to rescale the estimated effect of investment on household water rates without subsidies, following Equation (6). Each histogram reports the fraction of observations within each bin. Observations are jurisdiction-years.

centered at substantially higher values, around 6–8 percent. Thus, even before accounting for grants, subsidized lending materially reduces the annual debt-service burden associated with a given amount of borrowing.

Panel (B) reports the resulting no-subsidy cost multiplier, $\hat{C}_{jt}^{\text{unsub}}/\hat{C}_{jt}^{\text{base}}$, which is the adjustment applied to the estimated baseline water rate increase, following Equation (6). The mean predicted multiplier is 1.38, with a standard deviation of 0.37, and its median is 1.27. Thus, eliminating grants and subsidized loans increases the predicted annual financing cost by about 38 percent on average and by 27 percent for the median jurisdiction-year. The distribution is right-skewed, indicating that government support accounts for a much larger share of financing costs in some systems than in others. In the counterfactual analysis, we allow this multiplier to vary across jurisdiction-years rather than applying a common adjustment to all systems.

REFERENCES

- Allaire, Maura, Haowei Wu, and Upmanu Lall**, “National trends in drinking water quality violations,” *Proceedings of the National Academy of Sciences*, 2018, *115* (9), 2078–2083.
- Keiser, David A, Bhashkar Mazumder, David Molitor, and Joseph S Shapiro**, “Water Works: Causes and Consequences of Safe Drinking Water in America,” 2023.
- Price, James I and Matthew T Heberling**, “The effects of source water quality on drinking water treatment costs: a review and synthesis of empirical literature,” *Ecological economics*, 2018, *151*, 195–209.