INEFFICIENCIES IN LOCAL INFRASTRUCTURE: EVIDENCE FROM DRINKING WATER IN CALIFORNIA

PIERRE BODÉRÉ, GIULIA BRANCACCIO AND KARAM KANG

Abstract. Local elected officials oversee critical infrastructure, often with limited electoral accountability. This paper studies how this distorts investment across California drinking-water systems. We document that pollution responds to shifts in the political environment, motivating a dynamic model of public investment. We measure welfare returns to investment directly, combining event-study estimates of investment effectiveness and costs with those of residents' clean-water valuations, recovered from housing choices. Investment choices identify policymakers' responsiveness to those returns. Weak accountability diverts investment from high-return systems and conventional policies imperfectly offset these distortions: subsidy take-up is misaligned with returns, while uniform mandates induce welfare losses.

1. Introduction

Infrastructure investment in the United States has become a growing concern as spending struggles to keep pace with an aging capital stock. Local governments are overwhelmingly responsible, overseeing roughly half a trillion dollars annually in transportation and water infrastructure spending alone. Yet local officials making these investment decisions often face limited electoral accountability: local elections are low-turnout and often uncontested, weakening the residents' ability to reward or punish officials for their decisions. Weak electoral oversight may reduce officials' incentives to tailor investment to residents' needs, distorting investment decisions. Quantifying these distortions from observed infrastructure choices, however, presents a central identification challenge: low investment may either reflect low benefits to residents or officials' limited responsiveness to those benefits. We develop an empirical

---

Date: August 20, 2026.
We thank audiences at various seminars and conferences for helpful comments and suggestions. We are grateful to Arie Beresteanu, Ashley Langer, Lindsey Currier, and Nicolas Longuet-Marx for their insightful discussions. We thank Kartikeya Arora, Hunter DeClemente, Alejandro Herrera Caceido, Jiaoyang Li, and Genna Liu, and Jaime Valdes for excellent research assistance.
Bodéré (email: pierre.bodere@yale.edu): Yale School of Management.
Brancaccio (email: giulia.brancaccio@nyu.edu): New York University, CEPR, and NBER.
Kang (email: karam.kang@wisc.edu): University of Wisconsin-Madison and NBER.

framework that separates residents' benefits from the costs, constraints, and political incentives governing investment. We then quantify the misallocation in investment from limited electoral accountability, evaluate whether conventional infrastructure policies, such as expanded subsidies and direct investment mandates, can compensate for limited accountability.

We study drinking-water infrastructure in California, where hundreds of locally elected governments repeatedly make high-stakes capital-investment decisions. The setting combines substantial policy importance with unusually rich data. Safe drinking water remains a major policy challenge: despite substantial federal and state support, including $41.1 billion allocated since 1996 through Safe Drinking Water Act programs, roughly 10 percent of Americans are exposed to violations of drinking-water standards, with measurable health consequences (Allaire et al., 2018; Keiser et al., 2023). This challenge suggest that the problem may extend beyond fiscal capacity to how investment decisions are made. Drinking water also offers a particularly tractable setting relative to other public investments, such as public safety and education. Whereas outcomes in these settings depend heavily on factors beyond local officials' control, drinking-water quality is linked through a transparent production process to infrastructure investments that officials directly choose.

We assemble a panel of 520 California local governments—157 cities and 363 special districts serving more than 18 million residents—from 2003 to 2022. Because no comprehensive measure of drinking-water infrastructure investment exists, we construct investment events and their financing from municipal debt records, federal and state loans and grants, and intergovernmental transfers. We augment these records with detailed data on water quality and rates, local elections and voter turnout, government staffing, and housing transactions.

We first present evidence that the local political environment shapes water quality. Pollution is higher in jurisdictions where elections are less frequently contested and where local newspaper coverage is weaker, even after accounting for both persistent and time-varying determinants of investment costs and benefits. To move beyond cross-sectional comparisons, we exploit within-jurisdiction changes in election timing and show that aligning local elections with statewide contests, which changes the political environment in which local officials operate, without directly affecting investment fundamentals, leads to sustained reductions in pollution. An event study centered on the 2015 California Voter Participation Rights Act, which prompted many jurisdictions to move their elections on-cycle, provides further support for this pattern. Voters also appear to hold officials accountable where they can: among jurisdictions where incumbents typically face challengers, worse drinking-water quality increases the likelihood that an incumbent loses, even after instrumenting drinking-water pollution with ambient surface-water pollution.

Motivated by the role the political environment plays in infrastructure provision, we then develop and estimate a dynamic model of investment in public drinking-water systems. In each period, a locally elected governing body, the policymaker, chooses whether to undertake an infrastructure investment. Investment improves water quality but is costly and must be financed by water-rate increases paid by the residents. When an election is contested, residents decide whether to retain or replace the policymaker based, among other things, on current pollution and rates, which reflect past investment decisions. For the policymaker, investment can therefore improve re-election prospects when it benefits residents, but it also entails private costs of planning and implementation. The institutional features governing these electoral incentives, such as the likelihood of a contested election, the value of holding office, and the salience of water issues in voting, can be summarized by a single parameter governing the weight the policymaker places on residents' welfare relative to these private costs. We interpret this parameter as a measure of electoral accountability: when accountability is weaker, investment responds less to residents' net welfare gains from investment.

Identifying and estimating the model requires separating residents' preferences from policymakers' incentives. To this end, we estimate residents' net welfare gains from investment independently of observed investment choices and then use these gains to isolate the policymaker's electoral incentives. Accordingly, our estimation proceeds in multiple steps. First, we recover residents' willingness to pay for for drinking-water quality using housing transactions near water-system boundaries. Following the housing-choice framework of Bayer et al. (2007), we compare similar homes located near one another but served by different water systems and use buyer characteristics to identify heterogeneity in their willingness to pay. We strengthen identification with a new instrument for drinking-water quality based on ambient surface-water pollution. We then estimate the causal effects of investment on subsequent water quality and its cost using event-study designs. Combining these estimates allows us to construct the net welfare gains from investment in each jurisdiction and period.

We then use variation in these estimated welfare gains to identify the policymakers' payoff. Residents' benefits from investment vary across jurisdictions and over time as environmental conditions and water sources affect pollution, while financing conditions determine the costs residents bear. The responsiveness of observed investment to this variation identifies the weight that the policymaker places on residents' welfare, while systematic differences in investment propensity conditional on these benefits identify the policymaker's private cost of undertaking investment. We allow the weight on residents' welfare to vary with the institutional features governing electoral incentives, such as the likelihood of a contested election.

Our estimates show that residents place substantial value on drinking water quality. In our preferred specification, the average household is willing to pay approximately \$356 per year for a permanent one-standard-deviation reduction in drinking-water pollution. Infrastructure investment, in turn, produces sizable improvements in water quality: pollution falls by about 30 percent in years three through eight after investment, reaching a reduction of roughly 44 percent around year seven before the effect depreciates. Combining the preference and effectiveness estimates yields an average willingness to pay for investment of about \$250 per household per year.

The benefits and financial costs of investment vary substantially across households, jurisdictions, and over time. Although higher-income households have a greater willingness to pay for a given reduction in pollution, lower-income households tend to live in more polluted jurisdictions and therefore receive larger water-quality improvements from investment. As a result, willingness to pay for investment itself declines with income, ranging from roughly \$320 to \$115 per household per year across income groups. The financial cost borne by residents also varies with access to federal and state support. In jurisdictions predicted to rely more on federal and state financial support for investment, water rates remain essentially unchanged following investment; in other jurisdictions, by contrast, rates rise by about 10 percent on average over the subsequent decade.

Turning to policymakers, we find that the weight policymakers place on residents' welfare increases with the predicted probability of a contested election. Moving from the average probability of a contested election to certainty raises this weight by about 8 percent for cities and 46 percent for special districts. Administrative capacity also matters, particularly in systems with fewer skilled personnel capable of managing complex capital investments.

We use the estimated model to quantify the extent of misallocation in infrastructure investment arising from limited electoral accountability. To do so, we solve for the investment and pollution outcomes residents would face under consistently contested elections. We find that the distortion created by weak accountability is not simply too little aggregate investment. Instead, it is a failure to direct investment toward communities where its returns are greatest. Among systems in the most pollution-prone quintile, many of which are located in the San Joaquin Valley where persistent agricultural contamination coincides with weak local electoral competition, investment probabilities rise, leading to pollution reductions of approximately 6–9 percent of its average level. By contrast, investment declines modestly in cleaner jurisdictions where marginal benefits are lower. Investment also increases more in communities with lower income, homeownership, and college-educated shares, making the reallocation toward greater environmental need progressive along socioeconomic dimensions.

Finally, we ask whether conventional infrastructure policies can reproduce these gains while leaving the local political environment unchanged. Subsidies preserve local discretion and can, in principle, help officials select high-return projects, but this targeting advantage depends on voluntary take-up and on officials responding sufficiently to resident welfare. We find that existing federal and state subsidies lower aggregate pollution by less than 1% relative to a no-subsidy regime and by about 2% among the most pollution-prone systems. Doubling expected subsidy intensity increases these reductions to roughly 1% overall and 3–4% among high-pollution systems, still well below reduction generated by greater accountability. Uniform investment mandates eliminate the subsidies' take-up problem and generate larger aggregate pollution reductions than greater accountability, but at the cost of poor targeting: they require investment even in systems with clean water, high costs, or low resident willingness to pay, causing roughly one-third of residents to experience welfare losses. Thus, neither greater financial support nor uniform mandates fully substitute for institutions that align local policymakers' incentives with residents' welfare.

1.1. **Literature review.** A growing literature studies whether decentralized infrastructure decisions are efficient, in the context of schools (Cellini et al., 2010; Biasi et al., 2025), ports (Brancaccio et al., 2024), high-speed rail (Fajgelbaum et al., 2023), and wind farm siting (Kashner, 2026). Measuring efficiency requires quantifying both the benefits and the costs of investment. We estimate benefits using measures of residents' willingness to pay for water quality from housing choices, building on structural models of preferences for local amenities (Bayer et al., 2007, 2016; Barwick et al., 2024). Unlike Bayer et al. (2016) and Almagro and Domínguez-Iino (2025), where local amenities change in response to residential sorting, in our setting they evolve through infrastructure investment decisions made by local policymakers, whose responsiveness to residents depends on electoral incentives. We then combine household valuations with detailed data on drinking water quality, water charges, and local governments' financing of capital projects to estimate the welfare returns to local infrastructure investment.

We move beyond measuring inefficiency to investigating its causes, with a focus on electoral accountability. Existing research has highlighted several forces that can distort local public investment, including political favoritism (Burgess et al., 2015), cross-jurisdictional spillovers (Bordeu, 2023), the timing of infrastructure maintenance (Sileo, 2023), municipal financing frictions (Agrawal and Kim, 2022), and regulatory and pricing institutions (Timmins, 2002; Lim and Yurukoglu, 2018; Gowrisankaran et al., 2026).[1] Recent work also emphasizes dynamic commitment problems arising from the feedback between public infrastructure provision and private development and location choices (Hsiao, 2025; Ouazad and Kahn, 2025). Complementing this work, we study how policymakers' responsiveness to resident welfare and their capacity to plan and implement large capital projects shape infrastructure investment.

Our estimation strategy relates to work that uses observed policy choices to recover political incentives. Closest in spirit is Fajgelbaum et al. (2023), who combine a quantitative spatial model with referendum data to measure voters' responses to the local economic gains from California's high-speed rail and use the observed station design relative to alternatives to recover planners' political preferences. We first construct residents' welfare returns to infrastructure investment independently of observed policy choices and then use repeated investment choices to identify policymakers' responsiveness to these returns and their private costs of undertaking investment. Our setting features repeated, decentralized investment decisions by local governments, allowing responsiveness to vary with electoral contestedness and enabling us to quantify how weak electoral accountability distorts routine investment and the effectiveness of federal and state subsidies and mandates. This approach also complements studies that identify political incentives using exogenous changes in electoral institutions or constraints, including term limits (Besley and Case, 1995; List and Sturm, 2006; Sieg and Yoon, 2017, 2022), anti-corruption reforms (Avis et al., 2018), and campaign dynamics (Iaryczower et al., 2024).

[1]A complementary literature studies inefficiency in the allocation of water across agricultural users, emphasizing the role of legacy water rights and conveyance constraints (Rafey, 2023).

2. Drinking water infrastructure

This section describes the institutional environment governing drinking-water infrastructure investment in California. We first outline the basic operations of municipal water utilities. We then discuss how these investments are financed. Next, we describe the local governance structure and electoral institutions under which water utilities operate. Finally, we review the regulatory framework established by the Safe Drinking Water Act and the federal and state subsidy programs that support infrastructure investment. Together, these institutional features shape both the costs and political incentives associated with drinking-water infrastructure provision.

2.1. **Water utility operations.** Water utilities, typically run by local governments like cities, manage three critical operations. First, they source raw water by pumping it from surface water bodies like lakes and rivers, or from underground aquifers. Second, they process this water through treatment facilities to ensure it meets safety standards for human consumption. Third, they maintain storage systems and distribution networks to provide consumers with reliable, uninterrupted access to water. These operations depend on extensive infrastructure that demands ongoing financial investment and regular maintenance to function effectively.

2.2. **Financing water infrastructure.** To fund infrastructure investments, local governments rely on a combination of state and federal grants and debt financing, with the latter serviced by revenue from water sales. This financing structure reflects the organization of municipal water utilities as enterprises financially separate from general government operations. Unlike most government services, which are supported by tax revenue, water services are funded through user charges, and debt issued for water infrastructure is ultimately secured by these revenues (Posenau, 2021). As a result, new infrastructure investments typically necessitate higher water charges. This institutional feature has two important implications. First, from residents' perspective, infrastructure investment involves a genuine trade-off: it improves drinking water quality, but the cost is passed through directly to households in the form of higher water bills. Second, because investment is financed by ratepayers rather than out of the general budget, spending on water infrastructure does not crowd out other municipal services: a city that invests in its water system does not thereby forgo spending on its police force or schools.

In California, water rates face strict regulatory oversight. Proposition 218 empowers property owners to block water rate increases, while cost-of-service requirements limit price differentiation across households. These constraints result in simple pricing structures among water systems in our sample: 45% use uniform pricing (flat fixed and service charges), while 55% employ tiered systems with modest usage-based price variation. Even when water charges must increase to cover investments, regulatory requirements mandate that revenue increases be distributed proportionally across rate components based on customer service levels.[2]

2.3. **Local governance and electoral institutions.** In California, 84% of the population relies on local governments for their tap water. These providers fall into three categories: cities, dependent special districts, and independent special districts. Independent special districts are autonomous entities governed by their own boards, whose sole mandate is the provision of drinking water to residents. Dependent special districts, by contrast, are tied to a city or county, with governance exercised by the corresponding city council or county board of supervisors. Despite these institutional differences, all three types of providers share a common governance structure: they are overseen by multi-member governing bodies—city council, county board of supervisors, or independent special district board—whose members are directly elected by residents. Moreover, while professional general managers oversee technical and day-to-day operations, ultimate authority over capital investment and rate-setting rests with these elected bodies.[3]

Many elections for these local officials are uncontested or feature only a single candidate, and voter engagement in these contests is often low.[4] Election timing varies across jurisdictions and over time, with some elections held off-cycle (e.g., in odd-numbered years) and others consolidated with statewide or national elections. Institutional changes, including the California Voter Participation Act enacted in 2015 (SB 415), have led many jurisdictions to shift toward on-cycle elections. This law required jurisdictions with sufficiently low turnout in off-cycle elections to consolidate

[2]This is referred to as the proportionality requirement of property related fees, stipulated in Proposition 218 (see Article XIII D Section 6 of the California Constitution).

[3]Most cities in California operate under a council-manager system, in which the city council appoints a professional city manager to oversee day-to-day operations.

[4]These elections consistently see voter participation rates at only half the level of national elections, with even lower turnout when they are not synchronized with presidential or gubernatorial races (Hajnal and Lewis, 2003).

their elections with statewide contests.[5] In addition, some local governments had already chosen to shift the timing of their elections prior to the law. These cross-jurisdiction differences and over-time changes in election timing generate variation in local electoral environment that we leverage in Section 4.

2.4. **Water quality regulation and infrastructure subsidies.** The 1974 Safe Drinking Water Act (as amended in 1996) authorizes EPA to establish standards, generally maximum contaminant levels (MCLs), for more than 90 drinking-water contaminants. Public water systems must notify consumers within 24 hours of Tier 1 violations that pose acute health risks, including specified E. coli violations, and within 30 days for Tier 2 violations, which include most other MCL violations. The Act also requires systems to return to compliance, with regulatory responses escalating from compliance assistance and informal measures to administrative orders, penalties, and litigation when violations are serious, recurring, or unresolved.[6] Enforcement is not entirely automatic, however: states generally exercise primary enforcement responsibility, and EPA retains discretion over the timetable and mechanism for returning a system to compliance.

In addition to the regulatory standards for drinking water quality, the federal and state governments have provided substantial financial support to help local systems meet these requirements. Over the past 25 years, the Drinking Water State Revolving Fund (DWSRF) has served as the primary vehicle for this support, providing more than $40 billion in subsidized financing since its inception in 1997. This funding was expanded through the American Recovery and Reinvestment Act of 2009, and more recently, the Infrastructure Investment and Jobs Act of 2021.[7] In California, four statewide general obligation bond measures—Propositions 50 (2002), 84 (2006), 1 (2014), and 68 (2018)—have authorized billions of dollars for water quality and supply investments. More recently, California Senate Bill 200 (2019) created the Safe and Affordable Drinking Water Fund, which provides approximately $130 million annually to support safe drinking water projects.

By channeling large amounts of capital through low-interest loans and grants, these programs reduce the effective cost of drinking-water infrastructure investment,

[5] SB 415 applied when turnout in a jurisdiction's nonconcurrent elections was at least 25 percent below its average turnout in the previous four statewide general elections.

[6] Chronically noncompliant systems may ultimately be required to assess consolidation or transfer of ownership.

[7] The latter allocated approximately $50 billion for water infrastructure nationwide, including substantial additional funding for the DWSRF, although this expansion falls outside our study period.

implying that observed water-rate increases capture only a fraction of the true economic cost of providing safe drinking water. Importantly, the availability and intensity of such subsidies vary considerably across jurisdictions and over time, creating meaningful variation in the effective cost of infrastructure investment faced by local governments and their residents. For example, communities that are designated as "disadvantaged" may qualify for additional financial support, including zero-interest loans or principal forgiveness, and in California, the designation is based on the median household income in a water system's service area.

3. Data

We construct a comprehensive database of drinking water infrastructure investment and enrich it by integrating information on water quality, utility rates, municipal finances, local political conditions, housing market transactions, demographic and socioeconomic characteristics, and climate. This combined dataset enables a detailed analysis of both the determinants and consequences of infrastructure investment. Below, we describe our data sources and key variables.

3.1. **Data sources.** Our data construction relies primarily on California's Financial Transactions Reports (FTR), mandatory annual GAAP-compliant financial statements filed by local governments. These reports provide detailed information on revenues, expenditures, and capital outlays, as well as assets and liabilities, from which we extract water enterprise-specific measures. To more comprehensively capture infrastructure investment, we augment the FTR data with records from the California Debt and Investment Advisory Commission (CDIAC) and the EPA's State Revolving Fund (SRF) programs, the latter obtained through FOIA requests.

We measure drinking water quality using pollution data from the California Division of Drinking Water, which reports test results for more than 70 federally regulated contaminants at the public water system (PWS) level. We complement these data with records of Safe Drinking Water Act violations from the U.S. Environmental Protection Agency (EPA). Online Appendix OA1.2 provides details on the underlying data and construction of our water-quality measures. To link PWSs to their corresponding municipalities, we combine manual name matching, online verification, and institutional classification data from Dobbin et al. (2023).

In addition, we measure surface water quality around drinking water systems using the EPA Water Quality Portal (WQP, also known as STORET). These data provide contaminant readings from monitoring about 7,000 stations across streams, rivers, lakes, and reservoirs, including information on the chemical measured, sampling location, and date. To construct an aggregate surface-water pollution measure relevant to each jurisdiction's drinking-water risk, we use an attention-based neural network that learns how to weight nearby readings across pollutants, distances, and temporal lags, and we use its cross-fitted predictions of drinking-water MCL violations as the resulting pollution index.[8] Online Appendix OA1.3 provides a detailed description on the construction of this surface water quality measure.

To characterize the political environment shaping accountability, we focus on two dimensions: electoral competition and information availability. For competition, we identify whether elections for officials responsible for water investment are contested. City-level election data come from the California Election Data Archive (CEDA), while data for special districts are compiled from county election offices. For information, we measure local newspaper presence using the 3DLNews dataset (Ariyarathne and Nwala, 2024), which aggregates local news sources between 1996 and 2024 based on Google and Twitter scraping. We restrict attention to newspapers identified from Google search queries.

In addition, we measure the composition of the electorate at the jurisdiction level using individual-level voter registration data from L2, which covers the universe of registered voters and tracks turnout across federal, state, and local elections. The data include demographic and socioeconomic characteristics, such as income proxies.[9]

To measure local governments' administrative capacity, we use the California Government Compensation (GCC) database, which provides comprehensive information on public employees across all local governments, including job titles, compensation, and employment status.

To capture residents' preferences over water quality, we use CoreLogic housing transaction data, which we link to buyer characteristics from Home Mortgage Disclosure Act (HMDA) records and to tract-to-tract commuting patterns from Longitudinal Employer-Household Dynamics (LEHD) data. Finally, we incorporate demographic and socioeconomic characteristics from the American Community Survey,

[8]A fixed spatial radius or temporal window would be difficult to justify because we do not observe systems' precise intake locations, and the relationship between ambient readings and drinking-water quality depends on hydrologic transport, contaminant persistence, and source-water type. The attention-based approach instead allows the data to determine which monitoring stations, pollutants, distances, and lags are most informative for each system.

[9]Since the income reported by L2 is notoriously unreliable, we assign each voter the average income based on the block group of their address, and their gender and age.

mapped to municipalities using water service area boundaries from the California State Water Resources Control Board. We also include municipality-level precipitation measures from climate data sources.

3.2. **Scope of analysis.** We focus on cities and special districts that report water enterprise revenues and operate public water systems during the study period.[10] We further restrict the sample to entities primarily engaged in supplying water to local residents, rather than those focused on wholesale transactions or sales to business, industrial, or irrigation users.[11] We also exclude 25 local governments due to missing data (e.g., water rate information or parent government identifiers for dependent special districts) or because they serve fewer than 50 households.

Because our analysis concerns investments intended to improve water quality, particularly investments in treatment facilities, we further restrict the sample to systems that treat at least 50% of the water they supply rather than relying primarily on purchased water. The resulting sample consists of 520 local governments, including 157 cities (serving approximately 11.4 million residents) and 363 special districts (serving about 6.9 million residents). The majority of special districts (80%) are independent, while all dependent special districts in the sample are affiliated with county governments.

3.3. **Drinking water infrastructure investment.** Our novel dataset tracks California local governments' drinking water infrastructure investments, including investment frequency and funding sources. To our knowledge, this provides the first comprehensive view of local government infrastructure investment behavior on drinking water, expanding on prior work by Posenau (2021) and Sileo (2023).[12]

To identify infrastructure investment, we combine multiple sources of financing data (see Online Appendix OA1.1 for details). We first use indebtedness records from the FTR and cross-reference them with bond and loan issuance data from CDIAC to identify debt-financed investment. We then incorporate external funding sources

[10]Six county governments also meet these criteria but are excluded from the analysis.
[11]Specifically, we require that less than 25% of water revenues come from wholesale sales (including sales to other municipalities), less than 25% of retail water revenues come from sales to business, industrial, or irrigation users, and that average annual water sales exceed $15,000 in nominal terms.
[12]Posenau (2021) examines water utilities' debt contracts using our data sources, but doesn't investigate infrastructure investment choices. Sileo (2023) leverages Kentucky's Water Resource Information System (WRIS) to compile data on water systems and infrastructure projects, but does not study water service charges or the source of funding.

using intergovernmental transfers reported in FTR water enterprise income statements and Drinking Water State Revolving Fund (DWSRF) loan and grant records.[13] We aggregate all sources of financing—including debt, federal and state loans, and grants—to construct total investment funding at the jurisdiction-year level.[14] We define an investment event when total funding exceeds a minimum threshold of $500,000 (in 2020 dollars) to capture economically meaningful projects.[15] This approach relies on the premise that municipalities rarely finance infrastructure projects entirely from internal resources.[16]

Table 1 presents summary statistics separately for cities and special districts. Investment activity is relatively infrequent: in any given year, about 7% of cities and 5% of special districts undertake an investment. Among those that do invest, the median combined funding from debt and external sources is $10.3 million for cities and $5.7 million for special districts (in 2020 dollars). In contrast, mean values are substantially higher—$57.9 million and $39.6 million, respectively—indicating considerable right-skew in the distribution. A similar pattern appears in per-household terms: for special districts, the median investment of $2,223 per household is well below the mean of $9,464. Municipalities typically finance infrastructure through a combination of internal funds, debt, and external sources such as grants. We estimate that debt instruments—including municipal bonds and subsidized loans from federal and state governments—account for 62% of investment costs for cities and 56% for special districts.[17] As for the two types of debts, subsidized loans generally carry lower interest rates than municipal bonds: municipal bonds have an average coupon rate of 3.5% over our sample period, whereas 35% of observed DWSRF loans are interest-free and the stated interest rates on the remaining loans range from 1.1% in 2022 to 2.57% in 2010. Maturities, on the other hand, are similar: municipal bonds

[13]Specifically, FTR and CDIAC provide information on each debt instrument's principal amount, issuance year, stated purpose, pledged revenue sources, and any federal or state involvement. For external funding, we combine intergovernmental transfer records from FTR income statements with EPA DWSRF loan and grant data.

[14]Because financing may occur over multiple consecutive years, we group contiguous years with positive funding into investment spells, assign the first year of the spell as the investment year, and sum funding within the spell to obtain total investment.

[15]Our main results are robust to alternative thresholds.

[16]Using post-2017 FTR data for cities, which separately records equity-financed capital investment in water enterprises, we confirm this premise: the probability of an above-mean investment occurring without debt financing or external funding is only 1.3%.

[17]These figures are consistent with practitioner estimates of roughly 65% debt financing (Hansen and Mullin, 2022).

TABLE 1. Summary Statistics

	Cities	Special districts
A. Infrastructure investment		
Annual investment frequency	0.067	0.045
Conditional on any investment		
Amount debt-financed or externally funded		
Mean (in $ millions)	57.86	39.55
Median (in $ millions)	10.34	5.71
Per household, Mean (in $)	3,282	9,464
Per household, Median (in $)	1,818	2,223
% financed via debt (vs. grants)	62.02	55.99
B. Annual cost of water per household		
Mean (in $)	677.6 (427.7)	708.1 (448.9)
Median (in $)	553.0	579.6
C. Annual water pollution		
% readings above MCL	1.291	2.685
Any SDWA Tier I violations	0.005	0.007
Any SDWA Tier II violations	0.088	0.105
D. Political environment		
Fraction of contested elections, bi-annual	0.880	0.425
Fraction of on-cycle elections, bi-annual	0.890	0.620
Local newspaper presence (within 20 miles)	0.892	0.829
Local newspaper presence (within 10 miles)	0.752	0.576
Avg. number of skilled employees	42.53	10.23
Median net-equity-to-sales ratio	3.65	5.46
Number of governments	157	363

Notes: This table presents summary statistics of key variables used in our analysis, based on cities and special districts that treat water and provide drinking water services, for the period 2003–2022. The mean values are presented with standard deviations in parentheses, unless stated otherwise. All investment amounts and rates are adjusted to 2020 dollars using CPI. Conditional investment statistics are computed across government-year observations with positive investment and are based on the amount of debt financing and external funding in that year.

have an average maturity of 20.3 years, while the modal maturity of DWSRF loans is 20 years.

3.4. **Drinking water rates.** As described in Section 2.2, local governments finance water infrastructure investment through user fees and charges. Therefore, to assess the financial burden on water service users, we calculate the per-household cost by dividing the annual total of residential sales (from FTR water enterprise income statements) by the number of households in each municipality's water service area.[18] Panel B in Table 1 reveals a median annual cost of $553 per household for cities and $580 for special districts in 2020 dollars, with means exceeding these values in both cases. Notably, these costs show substantial variation across municipalities, as evidenced by large standard deviations.

3.5. **Drinking water quality.** Panel C of Table 1 summarizes drinking-water pollution in California water systems. Averaged across all jurisdiction-contaminant-year observations, the fraction of readings that exceed the applicable maximum contaminant level (MCL) is 1.3% for cities and 2.7% for special districts. A subset of MCL exceedances result in formal Safe Drinking Water Act (SDWA) violations. In a typical year, 8.8% of cities and 10.5% of special districts experience a Tier II violation. Tier I violations, which pose more acute health risks and require public notification within 24 hours, are rare, occurring in fewer than 1% of jurisdiction-year observations for both cities and special districts.

3.6. **Political environment.** Panel D of Table 1 reports the share of contested elections and the share of on-cycle elections for seats on the governing bodies of cities and special districts that provide drinking water, respectively.[19] Both margins – whether an election is contested and whether it coincides with a high-turnout statewide contest – shape the electoral pressure officials face, and thus the strength of the accountability mechanism at the center of our analysis.

There is substantial variation across entity types. Many local elections are uncontested, allowing incumbents to retain office by default, and this pattern is particularly pronounced among independent special districts (70% of elections, compared to 7% for dependent special districts and 12% for cities). For example, in Santa Fe Springs,

[18]Both FTR and the California Electronic Annual Reports (eAR) provide data on the number of households served by each public water system, but coverage is inconsistent. To address this, we instead estimate the number of households in each municipality using data from the American Community Survey combined with municipal shapefiles. We then use this measure to construct per-household residential water sales revenue. The eAR data also report average water rates for each public water system, though availability is similarly uneven. As a result, when our primary water rate measure falls outside a plausible range and eAR rate data are available, we substitute in the eAR-based rates. To validate our water rate measure, we directly contacted individual water utilities to obtain current rate schedules and documentation from rate hearing proceedings.

[19]We use mayor and city council elections for cities, county supervisor elections for dependent special districts, and water board elections for independent special districts. We identify the parent government of each dependent special district using FTR data. See Online Appendix OA1.4.1 and OA1.4.2 for how the share of contested elections and the share of on-cycle elections are constructed.

every city council election from 2003-2022 was contested, whereas the Ponderosa Community Services District held only two board elections over the same period.

Panel D of Table 1 also summarizes the extent of local newspaper presence across jurisdictions during the study period. For each year, we classify a newspaper organization as *active* if it publishes at least 10 articles in the dataset. We then measure local news presence by counting, for each jurisdiction-year, the number of active newspaper organizations headquartered within a 20-mile radius that overlaps with the jurisdiction's service area, as defined by the California State Water Resources Control Board.[20] Based on this measure, 15% of local governments are never covered by a nearby newspaper.

We assess governmental capacity to facilitate the planning and implementation of capital investments in drinking water infrastructure by two measures. The first is the ratio of net equity to sales revenue, calculated for the government's drinking water operations, which captures its financial capacity. The second is the number of highly compensated, skilled employees involved in drinking water. Panel D of Table 1 shows that the median net-equity-to-sales ratio is 3.65 for cities and 5.46 for special districts, and the average number of skilled employees is 42.5 for cities and 10.2 for special districts.

4. Suggestive evidence for infrastructure under-investment

This section shows that local water pollution is systematically associated with features of the political environment that both shape electoral incentives and are plausibly orthogonal to the underlying costs and benefits of infrastructure investment. These patterns persist when exploiting within-jurisdiction panel variation. Taken together, the results point to misallocation in the provision of safe drinking water and are consistent with a role for political accountability in shaping how public officials respond to constituents' needs.

4.1. **Cross-sectional evidence.** Our analysis is at the jurisdiction-pollutant-year level, indexed by j, p, and t. The main outcome, Pollution_{pjt}, is measured as the percentage of pollutant p readings that exceed its federal maximum contaminant level (MCL) across tests conducted in jurisdiction j during year t. Defining exceedance relative to the pollutant-specific MCL facilitates comparisons across contaminants. Our

[20] In practice, we use the service area of a public water system. If a jurisdiction is served by multiple systems, we take the maximum count of active newspaper organizations across those systems.

key explanatory variable, Political Env_j, captures features of the local political environment that shape electoral accountability—specifically, the fraction of contested elections and the presence of local newspapers, as discussed in Section 3.6. We construct these measures at the jurisdiction level by averaging over the full sample period. Because the pollution measure is often zero, we estimate a Poisson specification as follows:

$$\mathbb{E}[\text{Pollution}_{pjt}] = \exp\left(\beta \,\text{Political Env}_j \;+\; \mathbf{x}_{pjt}\alpha \;+\; \phi_{county(j)} \;+\; \mu_{pt}\right). \tag{1}$$

The specification includes county and period-by-pollutant fixed effects, $\phi_{county(j)}$ and μ_{pt}, respectively. In addition, we control for time-varying jurisdiction-level characteristics $\mathbf{x}_{jt}$, which includes demographics, such as income and the age distribution of residents; water system attributes; and precipitation. Finally, we control for the frequency of measurements across months for each pollutant-jurisdiction pair.

The first two columns of Table 2 show that jurisdictions with a higher frequency of contested elections and those with greater local newspaper presence exhibit significantly lower rates of pollution, even after controlling for factors that affect the underlying costs and benefits of infrastructure investment and water service. The magnitudes are economically meaningful. For example, the implied reduction in pollution associated with moving from no contested elections to fully contested elections is comparable to that associated with a one standard deviation increase in median household income (Column 1). Similarly, the presence of a nearby newspaper is associated with a sizable decline in pollution (Column 2). These patterns highlight the quantitative importance of the local political environment in shaping water quality outcomes.

4.2. **Exploiting panel-variation in election timing.** A causal interpretation of the patterns in Table 2 is limited because the local political environment can be correlated with underlying jurisdiction characteristics that also shape residents' preferences for drinking water quality and the costs of infrastructure investment. As an example, older residents are more exposed to the risk of water pollution, and are also more likely to participate in local politics. To move beyond cross-sectional comparisons, we exploit jurisdictions that change their election timing, primarily in response to institutional reforms such as the 2015 California Voter Participation Act (SB 415).

For our analysis, the key feature of election timing is that it alters the political environment facing local officials without directly changing the technological costs or

Table 2. Local Political Environment Matters for Water Quality

	Average % above MCL			
	(1)	(2)	(3)	(4)
Frequency of contested elections	-0.263*			
	(0.143)			
Any local newspaper within 20 miles		-0.330**		
		(0.151)		
Frequency of on-cycle elections			-0.334**	
			(0.153)	
Holding on-cycle elections				-0.132*
				(0.072)
Median household income (standardized)	-0.272***	-0.270***	-0.270***	-0.086
	(0.075)	(0.077)	(0.076)	(0.073)
Additional controls†	Yes	Yes	Yes	Yes
Pollutant-Year FE	Yes	Yes	Yes	Yes
County FE	Yes	Yes	Yes	No
Pollutant-Jurisdiction FE	No	No	No	Yes
Analysis-sample observations	251,432	251,432	251,432	251,432
Regression observations	48,454	48,454	48,454	7,525

Notes: This table reports Poisson regression estimates where the dependent variable is percentage of pollutant readings that exceed its federal maximum contaminant level (MCL) across tests at the pollutant-jurisdiction-year level, over the period 2003-2022. Standard errors, clustered at the government level, are shown in parentheses. Statistical significance is denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$. Analysis-sample observations count all pollutant-government-year observations with a nonmissing outcome in our sample, whereas regression observations refer to the observations actually used by the Poisson regressions after excluding observations dropped because of the fixed-effects structure. † Additional controls include population served, primary water source (surface vs. groundwater), socio-economic covariates (% white, population density, age distribution), the fraction of readings conducted in each calendar month, and precipitation measures (mean, 75^{th} percentile, 90^{th} percentile, and the fraction of days with no rain).

benefits of water infrastructure. On-cycle elections generate higher and more representative turnout (Hajnal and Lewis, 2003; Hajnal et al., 2022) and are considerably more likely to attract challengers (Appendix Table A2). To our knowledge, the relationship between election timing and candidate entry has not previously been documented. One possible explanation is that concurrent state and federal campaigns reduce the costs of entering local races by increasing voter attention, facilitating fundraising, and mobilizing parties and organized groups. On-cycle elections may also increase the expected return to entry by expanding and broadening the electorate. We discuss these mechanisms, as well as the potentially offsetting effect of incumbency advantage as documented by de Benedictis-Kessner (2018), in Appendix A.3. At the same time, it is worth emphasizing that for the analysis here, the specific mechanism through which election timing alters the political environment is not central.

Column (4) in Table 2 presents the estimate of the impact of holding on cycle election, controlling for a pollutant-by-jurisdiction and pollutant-by-year fixed effects, so that identification comes from changes in election timing within a jurisdiction. The estimated coefficient implies approximately 12 percent decline in the expected percentage of readings exceeding their MCL when elections move on-cycle.

We interpret the estimates in Column (4) as broad within-jurisdiction evidence, but the election-timing transitions have occurred across time, possibly for different institutional reasons. For this reason, in Appendix Figure A1 we conduct an event study by anchoring the identifying variation to a common institutional change: the 2015 reform of SB 415.[21,22] The results show no systematic differential trend before the reform and reveal that the largest impact happened in 2018 and 2019, three to four years after the law was enacted.

4.3. **Mechanism: Electoral Discipline.** We argue that electoral discipline is a key mechanism behind these results: when poor water quality increases the risk that incumbents lose office, officials facing greater electoral pressure have stronger incentives to prevent and reduce pollution.

This mechanism requires that drinking water quality matters to local voters. Several pieces of evidence suggest it does. Surveys document substantial public concern about drinking-water pollution, and California advocacy organizations such as the Community Water Center Action Fund and Clean Water Action California seek to raise the political salience of drinking-water issues through voter education, candidate endorsements, and grassroots campaigns. Local newspapers may further facilitate electoral response by informing residents about water quality conditions and utility performance.

Most directly, Table A1 in Appendix A.1 shows that, among jurisdictions for which the CEDA data allow us to distinguish incumbents from challengers, worse

[21]Because official records of election-timing changes are not systematically available, we identify regime changes using a Quandt Likelihood Ratio test for an unknown structural break (Andrews, 1993); Appendix A.2.1 provides details. We identify 59 jurisdictions that remained off-cycle, 24 that switched on-cycle before 2015, and 42 that switched after 2015.

[22]The exposed group consists of jurisdictions that moved from off-cycle to on-cycle elections in 2016, 2018, or 2020; the comparison group consists of jurisdictions with no post-2015 transition.

drinking-water quality increases the likelihood that at least one challenger wins a local election.[23] The relationship persists when pollution is instrumented with surface-water pollution. Thus, when incumbents face electoral competition, poor water quality carries consequences for their reelection prospects, consistent with voters being able to hold policymakers accountable for drinking water quality.

Taken together, the evidence supports electoral discipline as an important link between the local political environment and drinking water quality. In what follows, we develop and estimate a structural model to evaluate how local political environments shape the allocation of infrastructure investment and to examine its implications for policy.

5. Model

This section develops a dynamic model of investment in public drinking-water systems. Investment reduces pollution but raises water rates, and residents differ in how they weigh this trade-off. Investment decisions are made by an elected governing body (e.g., a city council or special district board). Residents therefore decide whether to retain or replace the members standing for election based on, among other things, current pollution and rates, which reflect the governing body's past investment decisions. These elections are the key disciplining force aligning the governing body's investment decision with residents' needs. Because the governing body comprises multiple members that face reelection on a staggered cycle, we model it as a single decision maker solving an infinite-horizon problem, hereafter referred to as the *policymaker*, and we treat seat-specific contests in a given period as a single election that either approves or disapproves the policymaker. Any change in the body's composition is thus treated as the replacement of the current policymaker.

At the start of each period t, the policymaker chooses whether to make a capital investment, $a_t = 1$, or not, $a_t = 0$. During the period, the realized pollution $q_t \in \mathbb{R}_+$ and the per-resident cost of service $r_t \in \mathbb{R}_+$ (covering capital, operations, and maintenance) are publicly observed.[24] At the end of the period, an election may occur, and the policymaker receives an office-holding benefit if she remains in office,

[23] Detailed candidate-level information is available only for the subset of elections covered by CEDA, primarily city and county elections, and not for the full set of special-district elections.

[24] Given (i) a balanced-budget requirement that user revenues cover per-period costs, (ii) institutional and legal constraints (see Section 2), and (iii) inelastic demand, we assume that all residents pay the same amount r_t for their water services.

whether by winning voter approval or by facing no challenger, and a payoff of zero otherwise.

5.1. **Resident preferences and voting.** Each resident i is characterized by income y_{it}. Let F_{yt} denote the cross-sectional distribution of income in the jurisdiction in period t. When an election occurs, resident i turns out with exogenous probability $\pi_t(y_{it})$; otherwise, she abstains. We parameterize resident i's trade-off between pollution and rates by $\omega(y_{it}) \in \mathbb{R}_+$, interpreted as her marginal willingness to pay to avoid pollution. The per-period utility from water service is

$$u(q_t, r_t, y_{it}) \;=\; -\omega(y_{it})\, q_t \;-\; r_t.$$

Conditional on turning out, resident i votes to approve the policymaker if her utility from water services exceeds her benchmark for the challenger,

$$u(q_t, r_t, y_{it}) \;\geq\; \bar{v} + \tilde{v}_{it},$$

where $\bar{v}$ captures the challenger's systematic appeal and $\tilde{v}_{it}$ is an idiosyncratic preference shock, reflecting other considerations that shape resident i's vote, such as candidate quality or other local issues. We assume that $\tilde{v}_{it}$ is uniformly distributed with density σ_v on a symmetric support, independent of (q_t, r_t, y_{it}). Note that σ_v captures the salience of water pollution in voting: when σ_v is small, a given change in water-related utility moves the choices of few voters, as may occur when voters have limited information about water quality or place greater weight on other local issues.

Finally, introduce aggregate uncertainty in the mapping from vote share to electoral victory. A popularity shock η_t shifts the realized vote share, where η_t is uniformly distributed with density σ_η on a symmetric support and independent of the idiosyncratic shocks. The policymaker wins if the expected vote share, shifted by the popularity shock, exceeds 1/2.

Define $\tilde{\pi}_t(y_{it}) \;\equiv\; \pi_t(y_{it}) \big/ \int \pi_t(y')\, dF_{yt}(y')$ as the turnout rate of residents with income (y_{it}), normalized by average turnout, so that $\tilde{\pi}_t(y_{it})$ reweighs the population income distribution to capture the income composition of voters who turn out. normalized by the population-average turnout rate. The policymaker's probability of approval is

$$\mu(q_t, r_t) \;=\; \frac{1}{2} \;+\; \sigma_\eta\, \sigma_v \left(\int \tilde{\pi}_t(y)\, u(q_t, r_t, y)\, dF_{yt}(y) \;-\; \bar{v} \right), \tag{2}$$

provided that the relevant arguments lie within the corresponding supports.

5.2. **Pollution and rate dynamics.** The endogenous component of drinking water pollution is governed by the investment stock, l_t^q, which evolves according to:

$$l_{t+1}^q \;=\; \tau_q a_t + (1-\chi_q) l_t^q, \tag{3}$$

where the parameter $\chi_q \in [0,1]$ denotes functional depreciation of the investment stock due to operational wear, and τ_q captures the effectiveness of investments at reducing pollution.[25] We assume that the jurisdiction's pollution is continuously distributed with mean:

$$\mathbb{E}[q_t] \;=\; \exp(\alpha_t - l_t^q), \tag{4}$$

where α_t captures the expected level of pollution in the absence of investment, fundamentals such as water source, local geography and soil, and climate conditions.

The per-period, per-resident cost of service reflects both operating expenditures and the recovery of past capital investments. The latter is captured by l_t^r, the per-resident stock of unrecovered investment costs. Its evolution depends on parameters τ_r and χ_r, capturing the per-resident cost of new investment a_t and the rate at which these costs are amortized over time.

$$l_{t+1}^r = \tau_r a_t + (1-\chi_r) l_t^r. \tag{5}$$

We assume that the per-resident water rate is continuously distributed with mean:

$$\mathbb{E}[r_t] = \gamma_t + l_t^r, \tag{6}$$

where γ_t captures the baseline, expected level of water rates, reflecting the jurisdictions' operating costs, in the absence of investment.

5.3. **The policymaker's investment decision.** When undertaking an investment, the policymaker incurs a private cost associated with the time and effort required to plan and implement the project. This cost cannot be passed on to residents and is

$$\kappa + \epsilon_t,$$

where ϵ_t is an i.i.d. logistic shock with unit scale.

The policymaker weighs this private cost against the future electoral benefit of investment, which depends on how residents value the resulting improvements in water quality relative to the associated rate increases.

[25] Note that infrastructure investments undertaken in period t become operational with a one-period lag. Thus, current pollution and rates reflect the stock of past investments, while a_t affects pollution and rates beginning in period $t+1$.

Let ρ_t denote the probability that an election is contested in period t, and let ν denote the benefit that the policymaker gets from being in office. Normalizing the payoff from not investing to zero, the policymaker's expected per-period payoff is therefore

$$\rho_t \nu \mu(q_t, r_t) - (\kappa + \epsilon_t) a. \tag{7}$$

Using the expression for the winning probability in equation (2), the electoral component of the policymaker's payoff can be written as

$$\rho_t \nu \left(\tfrac{1}{2} - \sigma_\eta \sigma_v \, \bar{v}\right) + \rho_t \nu \sigma_\eta \sigma_v \int \tilde{\pi}_t(y) u(q_t, r_t, y), dF_{yt}(y).$$

The first term is independent of the policymaker's action and therefore does not affect the investment decision. Collecting the institutional features governing the strength of electoral incentives into a single parameter,

$$\lambda_t \equiv \rho_t \nu \sigma_\eta \sigma_v. \tag{8}$$

we write the policymaker's decision-relevant per-period payoff as

$$\lambda_t \int \tilde{\pi}_t(y) u(q_t, r_t, y)\, dF_{yt}(y) - (\kappa + \epsilon_t) a. \tag{9}$$

We interpret λ_t as a reduced-form measure of electoral accountability, governing the extent to which resident welfare enters the policymaker's investment decision. It combines factors such as the probability of facing a challenger, the value of holding office, and the salience of water issues. A higher λ_t strengthens the alignment between policymaker incentives and constituent welfare, making investment more responsive to residents' welfare gains.

Let β denote the discount factor. Because the policymaker chooses the investment at the start of period t prior to the realization of water service outcomes, the public state is $\mathbf{s}_t \equiv (\alpha_t, \gamma_t, l_t^q, l_t^r)$. The expected value function can be rewritten as

$$\begin{aligned} V(\mathbf{s}_t) = \; \mathbb{E}_{\epsilon_t}\Bigg(\max_{a \in \{0,1\}} \Bigg\{ \lambda_t \int \tilde{\pi}_t(y)\, \mathbb{E}[u(q_t, r_t, y) \,|\, \mathbf{s}_t] \; dF_{yt}(y) \\ - \; (\kappa + \epsilon_t)\, a \; + \; \beta\, \mathbb{E}[V(\mathbf{s}_{t+1}) \,|\, a, \mathbf{s}_t] \Bigg\} \Bigg). \end{aligned} \tag{10}$$

Here, $\mathbb{E}_{\epsilon_t}(\cdot)$ averages over the contemporaneous shock ϵ_t; the inner expectations integrate over (q_t, r_t) conditional on $\mathbf{s}_t$ and over the next-period state $\mathbf{s}_{t+1}$ over $(a_t, \mathbf{s}_t)$, respectively.

6. Identification and estimation procedure

We develop the following empirical approach to separate residents' preferences from policymakers' incentives and recover the distortions induced by local discretion in infrastructure decisions. First, using our detailed data on drinking-water investment, pollution, and water rates, we estimate the cost and effectiveness of investment. Next, we estimate residents' marginal willingness to pay for pollution reduction from housing transactions, using an equilibrium model of housing choice across water-system boundaries following Bayer et al. (2007).[26] We combine these estimates to construct heterogeneous resident demand for investment, which we aggregate using individual-level turnout records to obtain the electorate's expected net benefits from investment. We then estimate the policymaker's payoff parameters using the rich variation in observed investment choices generated by repeated observations of hundreds of local governments over two decades. Below, we define the observables and structural primitives before describing each step of the estimation in detail.

6.1. **Observables and model primitives.** Jurisdictions (cities and special districts) are indexed by j. In each year t, the policymaker in jurisdiction j chooses an investment decision $a_{jt} \in \{0, 1\}$. Let $\boldsymbol{x}_{jt}$ collect exogenous state variables that affect water quality and costs—e.g., system size and water source—as well as its governance structure (city, dependent special district, or independent special district). Data provide the realized investment decisions a_{jt}, from which the investment stocks (l^q_{jt}, l^r_{jt}) are constructed, as well as the per-resident water rate r_{jt}, and covariates $\boldsymbol{x}_{jt}$.[27] Because water pollution is multidimensional, we summarize it using an annual jurisdiction-level index, q_{jt}. The index is the weighted sum of two standardized

[26]Two alternative approaches are possible but face important data limitations. The first would infer willingness to pay from avoidance behavior, such as purchases of bottled water or water filters. However, consumer-panel scanner data cannot provide the fine-grained geographic and time variation and individual demographics we get from housing transactions. The second would use voting behavior in water-infrastructure bond referenda or local elections. We identify only 133 relevant referenda between 1995 and 2022, and the associated bond measures often bundle water investments with other spending purposes. Local-election outcomes could, in principle, be especially informative and allow us to decompose λ further. However, the CEDA database does not cover most water-district boards; moreover, water-district elections are frequently uncontested.

[27]We initialize l^q_{jt} and l^r_{jt} using investment histories from 1990–2002. Because past investments depreciate at rates χ^q and χ^r, respectively, this finite pre-sample panel is sufficient to approximate the investment stocks at the start of the sample.

components: the number of SDWA violations and the number of distinct contaminants measured above their MCLs.[28] We also observe the jurisdiction-year-specific distribution of resident income y, denoted by F_{yjt}.

The structural primitives fall into three groups. First, the resident-side primitives are the marginal willingness to pay for water quality $\omega(y;\theta_\omega)$ and the income composition among voters who turnout, $\tilde{\pi}_{jt}(y;\theta_\pi)$, which we allow to be jurisdiction-year-specific. Second, the government-side primitives consist of the intertemporal discount factor, β, the policymaker's private cost of investment, and the weight placed on constituents' utility. We allow the latter two primitives to vary across jurisdictions and over time, parameterizing them as $\kappa(\boldsymbol{x}_{jt};\theta_\kappa)$ and $\lambda(\boldsymbol{x}_{jt};\theta_\lambda)$, respectively. We set $\beta = 0.92$, as the discount factor is not identified from choice-state data without additional restrictions (Rust, 1987). Finally, the state transition primitives consist of the parameters governing the effectiveness and durability of investment (τ_q, χ_q), the financial cost and amortization of investment for water rates $(\tau_{r,j}, \chi_r)$, and the baseline water quality and rate, α_{jt} and γ_{jt}, which we recover for each year and jurisdiction. We assume that the effect of investment on water pollution, τ_q, is common across jurisdictions, while allowing its effect on water rates, τ_r, to vary across jurisdictions, for example to reflect differences in federal and state financial support for investment.

6.2. **Effectiveness of investment on water pollution.** We begin by using a stacked event study approach to assess the impact of an investment across all investment events (Cengiz et al., 2019). For each investment event, e, we create an event-specific dataset, which includes observations from the treated government around the investment—between 3 years before and 10 years after the year of investment—and observations from all governments that did not have any major investments over the event window.[29] Since many readings are zero, we use a Poisson regression model with the conditional expectation as follows:

$$\mathbb{E}[q_{jte}|\{I^k_{jte}\}^{10}_{k=-3}, \boldsymbol{x}_{jt}] = \exp\left(\sum_{-3}^{10} \phi^q_k I^k_{jte} + \phi^q_x \boldsymbol{x}_{jt} + \psi^q_{je} + \mu^q_{te} + \eta^q_{jte}\right), \qquad (11)$$

where q_{jte} is the pollution measured at the drinking water facility of government j in year t, included in the e^{th} investment event dataset. The treatment dummy I^k_{jte} equals

[28]The two components capture different information. For some contaminants, such as coliform, violations depend on repeat-sampling procedures rather than an MCL comparison; even when an MCL applies, an exceedance does not necessarily constitute a violation.

[29]Our results are robust to perturbations to the time window and weighting of the control group observations depending on the propensity score.

1 if the investment event e occurred k years from year t by government j, and the specification includes event-specific government fixed effects and event-specific time effects, to ensure that identification comes from within an event.

We next use the event-study estimates to discipline the law of motion for pollution. We estimate (τ_q, χ_q) by fitting the geometric investment-response process in Equations (3) and (4) to the estimated event-study coefficients using nonlinear least squares. Here, τ_q governs the effect of investment on pollution and χ_q its persistence.

We also estimate α_{jt}, the expected level of pollution absent investment, by subtracting the cumulative effect of observed past investments implied by the fitted response $\hat{\phi}^q_k$ from expected pollution, $\widehat{\mathbb{E}}[q_{jt}]$, conditional on jurisdiction and year fixed effects and time-varying characteristics, such as weather.[30]

6.3. **Cost of investment.** To estimate how investment costs are passed on to residents through water charges (Section 2), we combine a stacked event-study design, similar to that used above, with additional structure to account for heterogeneity across governments and over time. Federal and state funding can substantially reduce the share of investment costs ultimately borne by residents. To capture this heterogeneity, we estimate the water-rate transition process separately for two groups defined by the predicted per-household amount of investment financed through debt. We construct this prediction using realized investment amounts and government characteristics considered by federal and state authorities in allocating funds, such as household income and population served.

We classify jurisdictions as high- or low-borrowing depending on whether their predicted per-household amount of debt-financed investment lies above or below the median. Letting $b \in L, H$ denote the borrowing group, we estimate the following stacked event-study specification separately for each group:

$$\log r_{jte} = \sum_{k=-3}^{10} \phi^r_{k,b} I^k_{jte} + \phi^r_{x,b} \boldsymbol{x}_{jt} + \psi^r_{je,b} + \mu^r_{te,b} + \eta^r_{jte}, \tag{12}$$

where r_{jte} is the per-resident water rate of government j in year t within the e^{th} event dataset, and $\phi^r_{x,b}\boldsymbol{x}_{jt}$ captures changes in ongoing maintenance and operating costs. We calibrate the persistence parameter χ_r using the typical maturity of investment-related debt in the data.

[30]Specifically, letting t_e denote the year of investment event e, we approximate the contribution of past investments to pollution in year t, $-l^q_{jt}$ in Equation (4), by $\sum_e \sum_{k=0}^{t-t_e} \hat{\phi}^q_k I^k_{jte}$.

As above, we recover γ_{jt} by estimating expected water rates conditional on jurisdiction and year fixed effects and time-varying characteristics and subtracting the cumulative effect of past investments implied by $\hat{\phi}^r_{k,b}$. We then measure $\tau_{r,j}$ by taking the average of $\hat{\phi}^r_{k,b}$ over the first five post-investment years and translating this log-rate effect into a level effect using $\hat{\gamma}_{jt}$.

6.4. **Resident preferences for water quality.** To recover the value placed by residents on pollution reduction, $\omega(\cdot;\theta_\omega)$, we rely on house transactions in California from 2003 to 2019. Consider a potential buyer i with income y_{it} choosing among houses h in a local housing market in year t. Houses vary in annualized transaction prices p_{ht}, physical and neighborhood characteristics $\boldsymbol{x}_{ht}$, the commuting distance to the buyer's workplace, d_{iht}. Each house is also located in a jurisdiction that supplies drinking water, denoted by $j(h)$, which determines pollution $q_{j(h)t}$ and the annual water rate $r_{j(h)t}$. Houses may additionally differ in an unobserved quality component ξ_{ht}, which may be correlated with price and drinking water pollution. The boundary-year fixed effects, $\psi_{b(h)t}$, are defined for houses located within 500 meters of water service jurisdiction boundary $b(h)$ in year t, and capture differences in amenities (other than the difference in pollution and rates) across jurisdictions. Because water system boundaries may not always lie entirely within a city or a school district, we additionally control for city-year and school district-year fixed effects respectively ($\zeta_{\text{city}(h)t}$ and $\iota_{\text{sch}(h)t}$), exploiting variation in water rates and pollution within jurisdiction, as illustrated in Appendix Figure B4.[31]

Buyer i's utility from purchasing house h in year t is

$$U_{iht} = \delta_{ht} + \theta_{qy}\log(y_{it})q_{j(h)t} + \theta_{py}\log(y_{it})(p_{ht} + r_{j(h)t}) + \theta_d d_{iht} + \varepsilon_{iht}, \tag{13}$$

where the idiosyncratic taste shocks ε_{iht} are assumed to follow the Type I extreme value distribution and the house-specific mean utilities δ_{ht} are specified as

$$\delta_{ht} = \boldsymbol{\theta}_x\boldsymbol{x}_{ht} + \theta_{q0}q_{j(h)t} + \theta_{p0}(p_{ht} + r_{j(h)t}) + \psi_{b(h)t} + \iota_{\text{sch}(h)t} + \zeta_{\text{city}(h)t} + \xi_{ht}.$$

This specification allows sensitivities to pollution and housing costs—the sum of annualized house price and water rate—to vary with buyer income. Let $\theta_\omega \equiv (\theta_{q0}, \theta_{qy}, \theta_{p0}, \theta_{py})$. Resident i's marginal willingness to pay for a reduction in pollution is

$$\omega(y_{it};\theta_\omega) = \frac{\theta_{q0} + \theta_{qy}\log(y_{it})}{\theta_{p0} + \theta_{py}\log(y_{it})}. \tag{14}$$

[31]On average, there are on average 2 public water systems run per city and 2 per school district. Appendix Figure B6 shows that housing characteristics do not change discretely at the boundaries.

Estimation of the house demand parameters in equation (13) then proceeds in two stages, following Bayer et al. (2007). In a first step, we estimate the house-specific mean utilities, denoted by δ_{ht}, along with the parameters governing buyer heterogeneity $(\theta_{qy}, \theta_{py}, \theta_d)$, via maximum likelihood. The estimation sample includes every residential property sold in California between 2003 and 2019. To minimize arbitrary restrictions on potential buyers' choice sets, we consider large geographic markets based on core-based statistical areas (CBSAs) and Appendix Figure B3 presents the resulting markets.

Our second step recovers the parameters governing residents' average willingness-to-pay by regressing the estimated mean utilities δ_{ht} on house characteristics, prices, and pollution. Due to the endogeneity of house prices, we instrument for prices using a variant of BLP instruments (Berry et al., 1995).[32] To strengthen identification beyond the boundary comparison, we instrument for drinking-water pollution using system-specific exposure to pollution in ambient surface waters, as detailed in Appendix OA1.3. The identifying assumption is that, among homes near the same water-system boundary and conditional on our fixed effects and controls, differential surface-water pollution exposure across systems affects house prices only through delivered drinking-water pollution.[33] Importantly, a natural concern with our instrument is that a polluted lake or river lowers house prices directly. The boundary design addresses this: houses on opposite sides of a boundary are equally exposed to the same nearby surface waters, so any direct disamenity is common to both sides and absorbed by the boundary-year fixed effect; what differs is only the serving system, and hence the quality of delivered water.

6.5. **Policymaker's parameters.** Observed investment choices identify the net value of investment to the policymaker (Rust, 1987; Hotz and Miller, 1993).[34] This net value combines the cumulative effect of investment on residents' expected welfare, weighted by electoral accountability λ, and the private cost of investment for the policymaker, κ. Investment choices alone, however, do not separately identify these components,

[32]Specifically, we capture the influence of nearby houses while excluding those in close proximity, exploiting variation in the exogenous characteristics of alternatives without picking up unobserved local amenities that are shared within small geographic areas.

[33]A similar design is used by Anderson (2020) to estimate the effects of long-run exposure to air pollution on mortality.

[34]We assume that the policymaker has rational expectations about the endogenous state variables (l^q, l^r), according to the transition processes estimated above. For tractability, although this could be relaxed, we assume that policymakers have perfect foresight over the evolution of exogenous state variables during our sample, and do not anticipate any change after 2022.

reflecting a well-known challenge in distinguishing policymakers' own incentives from residents' preferences using observed policy choices (Mian et al., 2010; Fajgelbaum et al., 2023). To address this challenge, our identification strategy exploits the fact that every component of residents' aggregate net benefit of investment can be measured directly. Resident preferences are estimated from housing transactions in the previous step, and turnout rates $\tilde{\pi}_t(y)$ can be estimated directly from individual-level voting records.[35]

Given these objects and the estimated state transitions, separating κ and λ relies on variation in the effect of investment on residents' expected net welfare that does not directly affect the policymaker's private cost. We obtain such variation from changes in both the benefits and the financial costs of investment to residents. The benefits vary with pollution, which changes over time with environmental conditions, while residents' expected financial burden varies with financing conditions, including the likelihood of receiving federal and state support. The responsiveness of observed investment to these shifts identifies λ. Conditional on this responsiveness, κ is identified from systematic differences in investment that are not explained by variation in residents' net benefits.

We parameterize λ and κ parsimoniously. In principle, λ_{jt} could depend flexibly on $\boldsymbol{x}_{jt}$. Such a specification, however, would require substantial independent variation in residents' net benefits from investment to separately identify how each covariate affects electoral accountability. We therefore allow λ_{jt} to depend on a single index capturing the strength of electoral discipline: the predicted probability of a contested election, P^C_{jt}, motivated by the evidence in Table 2. Among the many institutional features that may shape policymakers' responsiveness to residents, we focus on electoral contestedness, an important dimension of electoral discipline that provides observable variation in the strength of political accountability. We estimate P^C_{jt} using a Probit model with jurisdiction and year fixed effects and a time-varying indicator for whether elections are held on-cycle. Because this Probit model is dominated by jurisdiction fixed effects, P^C_{jt} primarily captures persistent institutional features that shape electoral discipline—such as the local pool of potential challengers or civic engagement—while also allowing electoral incentives to vary with election timing. Using predicted rather than realized contestedness further avoids conditioning on an

[35]Specifically, to estimate this object, we divide residents into three income groups and, for each jurisdiction and two-year period, predict the fraction of voters in each income bin, conditional on turning out, from a multinomial Logit estimated on data combining individual turnout histories with income estimates.

outcome which may itself respond to the policymaker's investment choices. With that, we specify λ_{jt} as

$$\lambda_{jt} = \exp\left(\theta_{\lambda_0} + \theta_{\lambda_1} P^C_{jt}\right), \qquad P^C_{jt} = \Pr\left(\text{Contested}_{jt} = 1 \mid \boldsymbol{x}_{jt}; \theta_{\lambda_2}\right),$$

where θ_{λ_0} governs the baseline level of responsiveness to resident welfare, while θ_{λ_1} captures how this responsiveness varies with predicted electoral contestedness.

Finally, we specify the private investment cost κ_t as a function of factors that capture the government's capacity to plan and implement capital projects. Specifically, κ_t depends linearly on the log average number of employees dedicated to water-utility operations over the sample period, and an indicator for systems with government net-equity-to-sales ratio below its 25th percentile, which captures financial health. We also include county fixed effects to absorb persistent regional differences in implementation and planning costs, such as unobserved urban–rural differences in access to engineers and contractors.

Given these parameterizations, we estimate the turnout rates from the turnout data, and the contested-election parameters, θ_{λ_2}, from the election data. We then jointly estimate the remaining parameters, $(\theta_{\lambda_0}, \theta_{\lambda_1}, \theta_\kappa)$, by maximum likelihood, solving the policymaker's dynamic problem at each parameter guess using the nested fixed-point algorithm of Rust (1987). See Online Appendix OA3 for details.

7. Estimation results

This section reports the model estimates. We begin by the estimated pollution and cost dynamics, highlighting the effectiveness of capital investments in reducing drinking water pollution and the heterogeneous costs borne by residents of different jurisdictions. We then turn to residents' preferences, summarized by their marginal willingness to pay for pollution reduction, $\omega(y)$. Finally, we present estimates of the policymaker's private cost in planning and undertaking an investment and the extent to which their investment decisions are responsive to resident welfare.

7.1. **Effectiveness and cost of investments.** Panel (A) of Figure 1 plots the event-study coefficients from (11). Relative to the pre-investment baseline, water-system upgrades produce a sharp and statistically significant decline in contamination: in years 3–8 after investment drinking water pollution, which combines SWDA violations and contaminant readings above the MCL as described in Section 6.1, falls by around 30% and up to 44% 7 years after investment. This magnitude closely matches the 0.5

Figure 1. Effects of Investment on Water Pollution and Rates

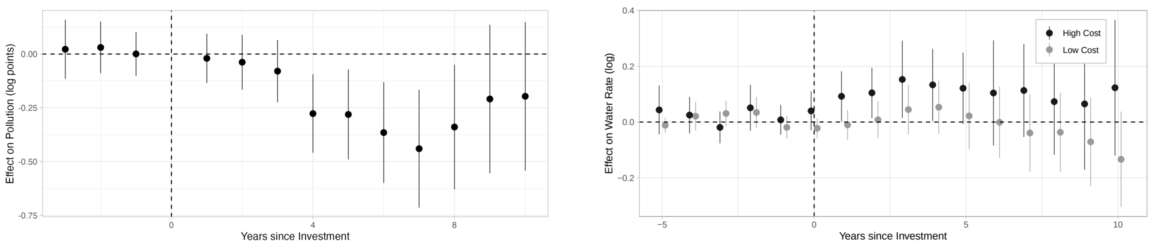

(A) Water pollution (B) Water rates

Notes: The figure reports estimated coefficients for the dynamic effects of local government drinking water infrastructure investment on water pollution (Panel A) and average per-household water rates (Panel B), using data over the period 2003-2022, based on Equations (11) and (12) respectively. Water pollution q_{jt} is an index constructed from the number of SDWA violations and the number of distinct contaminants above their maximum contaminant levels (MCLs); see Section 6.1. For water rates, we estimate separate effects for jurisdictions with high versus low predicted per-household amount of debt-financed investments. These groups are labeled "High Cost" and "Low Cost," respectively, in Panel B.

pp drop (roughly 33 %) reported by Keiser et al. (2023) for systems receiving Drinking Water State Revolving Fund loans. Online Appendix Figure OA4 shows that the effect is fairly uniform across systems with different observable characteristics. The improvement erodes over time and vanishes in ten years. This maps to $\hat{\chi}_q = 0.27$ and $\hat{\tau}_q = 2.24$.

Panel (B) of Figure 1 presents the event study coefficients from (12), separately for the high- and low-borrowing jurisdictions. Water rates remain largely unchanged following investment in low-borrowing jurisdictions. In contrast, high-borrowing jurisdictions experience a persistent increase in rates, averaging roughly 10% over the ten years following investment.

7.2. **Resident preferences for water quality.** Table 3 presents the demand estimates specified in (13). The top panel reports the coefficients governing heterogeneity in household preferences, and the bottom panel reports coefficients controlling common preferences. Most importantly, we estimate a negative and significant coefficient on drinking water pollution, implying that households value cleaner drinking water. This result is consistent with the evidence in Appendix Figure B2, which shows that house prices within jurisdiction fall significantly following a Tier I drinking water violation. Translating the estimates of Table 3 into willingness to pay, we find that

Table 3. Residents Value Drinking Water Quality Heterogeneously

First Step: Heterogeneous preferences			
Income (log) × Drinking-water pollution	0.012 (0.002)		
Income (log) × Annualized housing cost	34.566 (0.258)		
Income (log) × Block-group income	0.515 (0.049)		
Commute distance	-6.769 (0.151)		
Negative log-likelihood	35,573,410.628		
Second Step: Common preferences (dependent variable: $\hat{\delta}_{ht}$)			
	(1)	(2)	(3)
Annualized housing cost	-62.34 (1.052)	-62.56 (1.098)	-62.25 (1.155)
Drinking water pollution	-0.018 (0.002)	-0.036 (0.003)	-0.022 (0.006)
College educated share	1.059 (0.048)	1.049 (0.050)	0.972 (0.048)
Number of bedrooms (log)	0.029 (0.007)	0.031 (0.008)	0.034 (0.007)
Other house and neighborhood char.	Yes	Yes	Yes
Boundary × year FE and market FE	Yes	Yes	Yes
City × year FE	No	No	Yes
School × year FE	No	No	Yes
Drinking-water pollution IV	No	Yes	Yes
Observations	977,197	906,183	906,183
First-stage Wald stat: House price	2,508.5	2,256.4	1,997.7
First-stage Wald stat: Pollution		258.8	203.2

Notes: The first-step panel reports maximum-likelihood parameter estimates, with standard errors in parentheses. Income (log) stands for centered and standardized log of buyer income. The second-step panel reports IV regressions of estimated mean utility, $\hat{\delta}_{ht}$, on housing costs, drinking-water pollution, house characteristics (bedrooms, baths, building and lot square footage, year-built bins, land use, and pool) and neighborhood demographics (race/ethnicity shares, college share, and median income). In all specifications, we instrument for the house-price component of annualized housing cost, using a variant of BLP instruments based on the house's own exogenous characteristics and the exogenous characteristics of alternative houses located 1-5 kilometers away. Column (1) treats drinking-water pollution as exogenous, whereas Columns (2) and (3) instrument for it using the system-specific surface-water pollution index constructed from nearby ambient surface-water readings. Standard errors and Wald statistics are clustered at the public-water-system-by-year level.

the average household is willing to pay $356 for a one-standard-deviation *permanent* reduction in drinking-water pollution.[36]

While residents generally value cleaner drinking water, the magnitude of this valuation varies systematically with income. Higher-income residents exhibit a higher

[36]The permanence assumption facilitates interpretation of the willingness-to-pay measure. In practice, the effect of infrastructure investment on pollution eventually reverts toward its pre-investment level, as shown in Figure 1.

FIGURE 2. Households Willingness to Pay for Investment

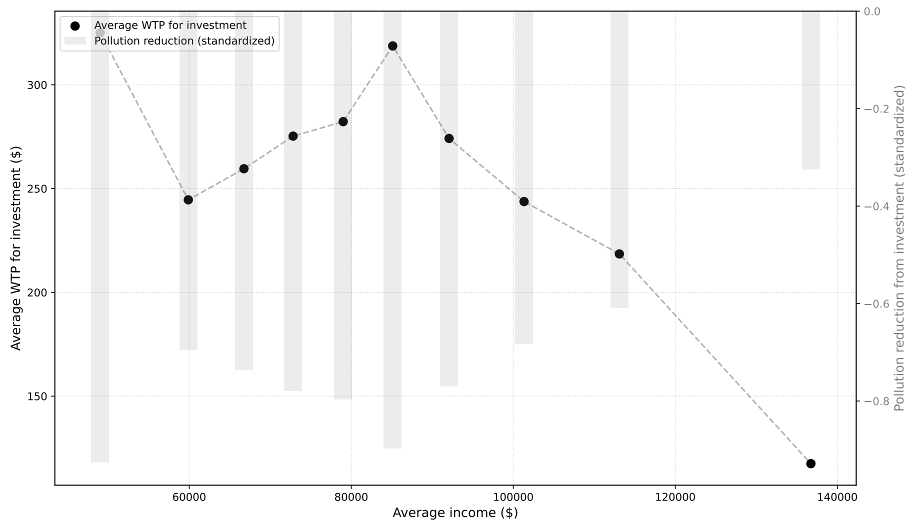

Notes: This figure shows average willingness to pay for drinking-water infrastructure investment across income bins, computed as the product of two terms: $\text{wtp}^{INV}(y) = \omega(y) \times \Delta q(y)$. Here, $\omega(y)$ captures heterogeneity in willingness to pay for pollution reduction by household income y, and $\Delta q(y)$ denotes the average reduction in pollution associated with investment in the jurisdictions where households with income y live. The dashed line and black dots, read against the left axis, plot $\text{wtp}^{INV}(y)$ in CPI-adjusted 2020 dollars. The gray bars, read against the right axis, show $\Delta q(y)$, measured in standard deviations of pollution across jurisdiction-years.

marginal willingness to pay for reducing in drinking water pollution, as shown in Appendix Figure B7. This pattern reflects two offsetting forces presented in the top panel of Table 3. First, consistent with the existing literature (Bayer et al., 2007, 2016), higher-income buyers are less sensitive to housing costs, which raises their willingness to pay for a given reduction in water pollution. Second, higher-income buyers are also less sensitive to drinking water pollution, as indicated by the positive and significant interaction between buyer income and pollution. One possible explanation is that higher-income households have better access to avoidance technologies, such as home filtration systems, which reduce their effective exposure to contaminated public water and therefore attenuate their revealed preference for reduction in system-level water pollution. Quantitatively, however, the weaker sensitivity to housing costs dominates the weaker sensitivity to pollution, resulting in a marginal willingness to pay for water quality that rises with income.

Figure 2 translates these preference estimates into households' annual willingness to pay for the pollution reduction induced by an infrastructure investment, assuming for illustrative purposes that this reduction is permanent.[37] Despite the fact that willingness to pay for a given improvement in water quality rises with income, willingness to pay for investment itself declines with income. The reason is that the benefits of investment are not the same across jurisdictions. Combining the estimated effect of investment on pollution from Figure 1 with households' baseline pollution exposure, we find that investment produces larger pollution reductions for lower-income households, who tend to live in more polluted jurisdictions. These gains are shown by the gray bars in Figure 2. Multiplying these investment-induced improvements in water quality by households' willingness to pay for pollution reductions yields the black dots: overall willingness to pay for drinking-water infrastructure investment decreases with income, ranging from roughly \$115 to \$320 across income groups. These magnitudes are substantial relative to an average annual household water bill of about \$700 (Table 1), but remain below typical expenditures on bottled water as an avoidance strategy. U.S. bottled-water consumption averaged 47.3 gallons per person in 2024; at an average retail expenditure of \$3.08 per gallon, this implies annual expenditures of approximately \$580 for a family of four.

Our revealed-preference estimates should be interpreted as households' perceived benefits of cleaner drinking water rather than as a complete measure of its social benefits. Households may not fully internalize the health consequences of pollution: eliminating readings above health standards corresponds to a 0.58-standard-deviation reduction in our pollution index, which households value at about \$206 per household-year, whereas Keiser et al. (2023) imply health benefits of roughly \$700 per person-year. Although the estimates are not directly comparable, the gap suggests that revealed willingness to pay may substantially understate health benefits. Our demand model also abstracts from dynamic considerations; as Bishop and Murphy (2019) show, mean reversion in water pollution can further bias such estimates downward. Thus, our estimates are best viewed as a conservative measure of the welfare gains from pollution reduction.

7.3. **How local political factors shape infrastructure investment.** Table 4 presents the estimated parameters governing policymakers' private investment costs, κ, and the degree of electoral accountability, λ. Consistent with the evidence in Section 4, the coefficient on the predicted probability of holding a contested election is positive and statistically significant. Moving from the average probability of a

[37]In estimation, we take into account the depreciation of investment benefits over time as estimated in Figure 1.

Table 4. Estimates of the Role of Political Factors in Investment

	Coefficient
κ: Private investment cost	
Number of skilled employees (log)	-0.192 (0.035)
Low government net-equity-to-sales ratio	0.007 (0.105)
County FE	Yes
λ: Accountability	
Constant	-8.528 (0.378)
Probability of holding a contested elections	0.695 (0.341)
Observations	6,175
Avg. Neg. log-likelihood	0.211

Notes: This table presents coefficient estimates for $(\theta_\kappa, \theta_\lambda)$ and their standard errors in parentheses. Coefficients are estimated using Nested Fixed Point (NFXP) maximum likelihood and standard errors are computed using the approximate Hessian at the parameter estimates.

contested election (Table 1) to certainty would increase λ by about 8% for cities and 46% for special districts.

We also find that the private cost of investment, κ, decreases with the number of highly compensated, skilled employees in drinking water. Ignoring the effect that changing private costs would have on continuation values, increasing the skilled workforce by 10% would increase the log-odds of investment by around 2%. This result echoes recent findings on the role of skilled employees in shaping the ability of local governments to invest in infrastructure (Liscow et al., 2025).

8. Quantification and policy implications of accountability

Our results in Section 4 show that weaker electoral discipline is associated with worse drinking-water quality, suggesting that limited political accountability may distort the allocation of infrastructure. The structural estimates in Table 4 reinforce this interpretation: the weight that policymakers place on resident welfare, λ, is substantially lower in jurisdictions where contested elections are less frequent. We use the estimated model to quantify the resulting distortions and to ask whether conventional infrastructure policies can offset them.

We first consider a high-accountability counterfactual in which every jurisdiction is assigned the degree of electoral accountability λ associated with consistently contested elections, while holding all other model primitives and the existing subsidy

Figure 3. Accountability and the Responsiveness of Investment

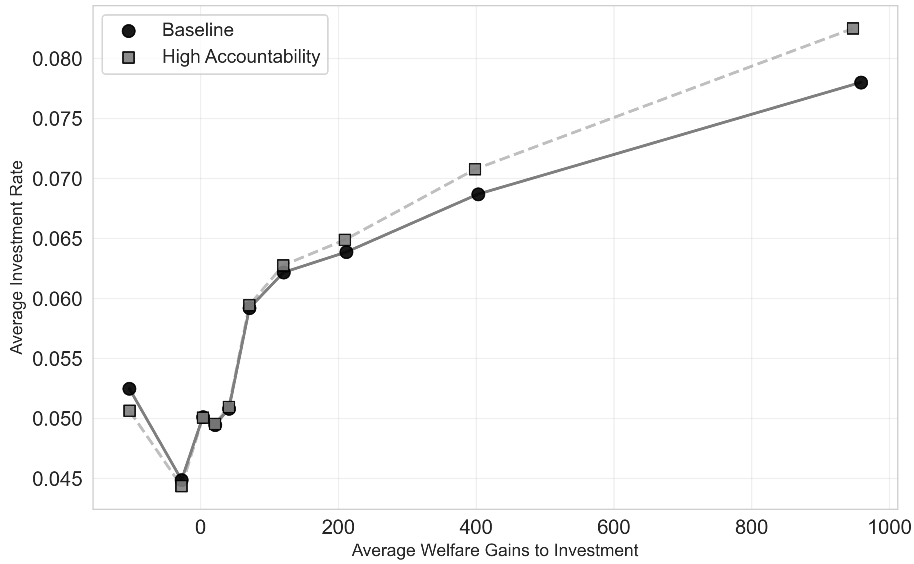

Notes: This figure plots the average investment probability as a function of the net welfare gains from investment under the baseline and high-accountability regimes. In the high-accountability counterfactual, each jurisdiction's accountability index is set to the level associated with consistently contested elections.

environment fixed.[38] This exercise quantifies how greater accountability changes the level and allocation of investment, and identifies where limited accountability generates the largest investment shortfalls. We then ask whether traditional policies that target investment, such as expanded subsidies or investment mandates, can reproduce these gains while leaving accountability at its baseline level.

8.1. **Accountability and the allocation of infrastructure investment.** To quantify the extent of infrastructure investment misallocation associated with limited political accountability, we compare outcomes under the estimated baseline regime with those under the high-accountability counterfactual. Specifically, we initialize each system at its observed state in 2003, solve the policymaker's dynamic investment problem under each regime, and simulate 1,000 paths of investment, pollution, and water rates through 2022.

[38]Although this counterfactual represents an ambitious upper-bound benchmark rather than a literal policy proposal, institutional reforms can meaningfully shift electoral competition. For example, the 2015 California Voter Participation Rights Act moved many local elections on-cycle, and our evidence indicates that on-cycle elections are substantially more likely to be contested (Appendix Table A2.

The key mechanism in this exercise follows directly from the estimated relationship between electoral accountability and policymakers' responsiveness to resident welfare. Table 4 shows that the welfare weight λ increases with a jurisdiction's probability of holding contested elections. Greater accountability therefore makes investment decisions more responsive to the net welfare gains generated by infrastructure improvements. Figure 3 illustrates this mechanism by plotting average investment probabilities against the net welfare gains from investment under the baseline and high-accountability regimes. Investment probabilities generally increase with welfare gains in both regimes. The two investment schedules are very similar at low welfare gains but diverge as the gains increase. In the highest welfare-gain group, the investment probability is approximately 0.7 percentage points (or equivalently 7% of the baseline probability) higher under high accountability. Thus, the investment shortfalls associated with limited accountability are concentrated where the welfare gains from investment are greatest.

Figure 4(A) shows the geographic distribution of the resulting investment changes. The largest increases are concentrated in the San Joaquin Valley, an agricultural region in Central California with persistent water-quality problems (Balazs et al., 2011). In some jurisdictions, investment rises by as much as 15 percent relative to the average level of investment in the baseline. Many of these communities also have weak electoral competition, limiting local officials' incentives to address serious water-quality problems even when funding is available.[39]

The additional investment translates into water-quality gains precisely where the underlying pollution is greatest. Figure 4(B) reports the difference in pollution between the high-accountability counterfactual and the baseline for systems in the bottom and top quintiles of underlying pollution, measured by each jurisdiction's average α_{jt} defined in Section 6.2. Pollution remains essentially unchanged among systems in the lowest-pollution quintile but declines by 8% of the average level of pollution among systems in the highest-pollution quintile.

A similar pattern emerges along socioeconomic dimensions: Online Appendix Figure OA6 shows larger increases in investment under the high-accountability regime among systems serving communities with lower college-educated shares, homeownership rates, and median household income. Thus, stronger accountability directs

[39] A 2018 report by the nonprofit Community Water Center observes that "In the local water boards of the southern San Joaquin Valley, contested elections are the exception, not the rule."

FIGURE 4. Heterogeneous Effects of Greater Accountability

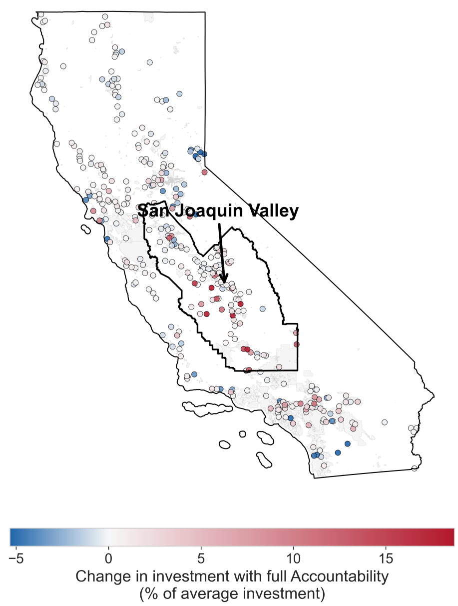

(A) % Change in Investment Probability

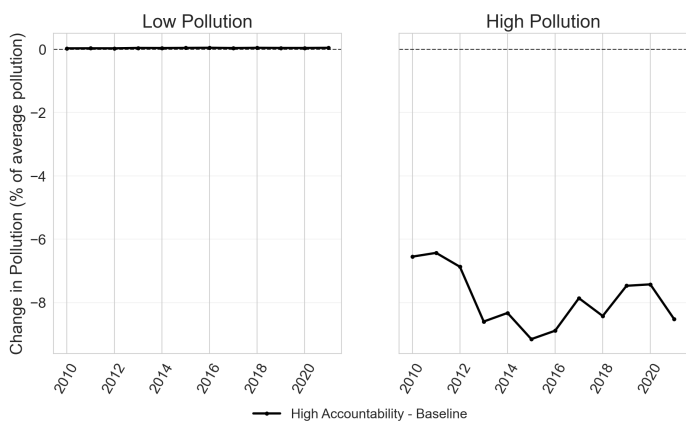

(B) Change in Pollution by Underlying Pollution

Notes: Panel (A) maps, for each jurisdiction, the percentage change in the average investment probability under the high-accountability counterfactual relative to the baseline, averaged across 1,000 simulated paths over 2003–2022. Panel (B) plots the annual difference in pollution between the high-accountability counterfactual and the baseline, expressed as a percentage of the cross-jurisdiction average of pollution. Jurisdictions are grouped according to their average underlying pollution, defined as the sample mean of α_{jt} defined in Section 6.2; the left and right panels report results for the bottom and top quintiles, respectively.

investment toward communities facing both greater environmental need and larger socioeconomic disadvantage.

8.2. **Can infrastructure policies offset the effects of weak accountability?** Section 8.1 shows that weak accountability creates the largest investment shortfalls where infrastructure improvements are most valuable; we now ask whether policies that directly alter investment incentives can offset these distortions. We consider two prominent approaches—expanded subsidies and investment mandates—which differ in the discretion they leave to local governments. A uniform mandate requires systems to invest according to a common schedule and cannot readily accommodate heterogeneity in water quality, investment costs, and resident preferences. Subsidies instead preserve local discretion allowing governments, in principle, to direct investment toward projects with the greatest local returns (Oates, 1972); this advantage, however, depends on local investment choices being sufficiently responsive to resident welfare.

In both counterfactuals, we hold λ and κ fixed at their baseline level, thereby asking how far conventional infrastructure policies can mitigate the consequences of weak accountability without changing electoral institutions.

8.2.1. *Expanding subsidies.* We consider two counterfactual subsidy regimes. First, to measure the contribution of existing financial support, we eliminate federal and state grants and replace subsidized loans with market-rate borrowing. Second, we consider an expansion that doubles each jurisdiction's expected subsidy intensity. Under each regime, we recompute the increase in water rates required to finance an investment, accounting for the share covered by grants, subsidized loans, and locally financed debt.[40] We then resolve local officials' dynamic investment problems and simulate the resulting paths of investment, pollution, and water rates.

Figure 5 shows that existing subsidies produce only modest improvements in drinking water quality. Relative to the no-subsidy counterfactual, the baseline subsidy regime lowers aggregate pollution by less than 1% of its average level and by approximately 2% among systems in the highest quintile of underlying pollution. These modest effects reflect the limited response of investment to subsidies, rather

[40]Online Appendix OA4.2 details the construction of counterfactual water-rate increases. By our calculations, removing grants and subsidized loans increase the annual financing cost of investment by 38% on average and 27% at the median.

FIGURE 5. Pollution Effects of Drinking-Water Subsidies

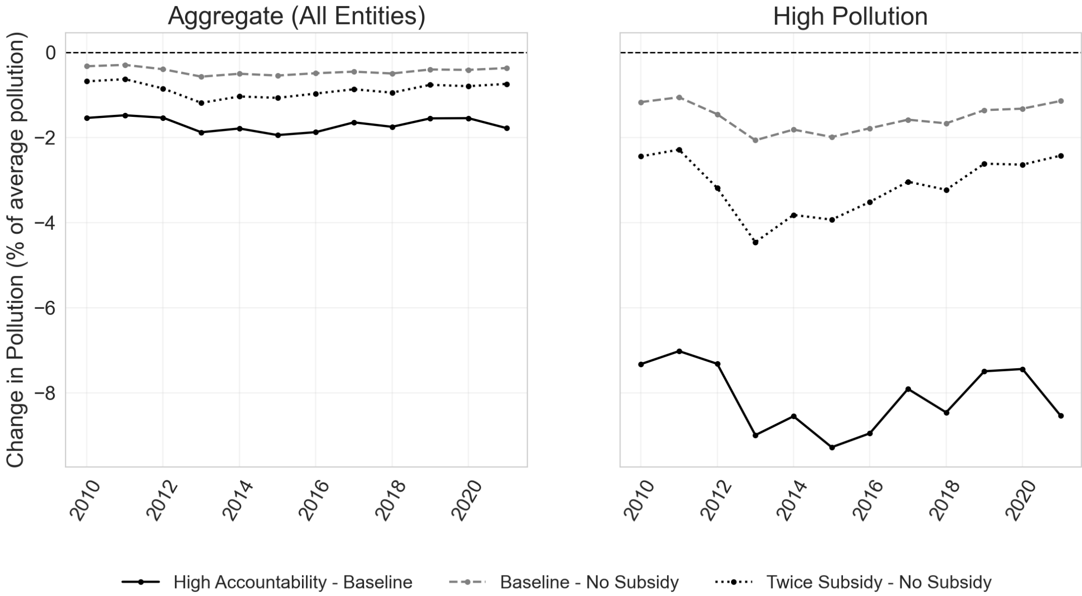

Notes: The figure reports annual differences in pollution across counterfactual regimes, averaged across simulated paths and expressed as a percentage of the cross-jurisdiction average of pollution. The left panel includes all jurisdictions, while the right panel is restricted to jurisdictions in the highest quintile of average underlying pollution, defined using the sample mean of α_{jt}. "Baseline – No Subsidy" reports pollution under the baseline subsidy regime relative to the no-subsidy counterfactual, which removes grants and subsidized loans, and therefore measures the contribution of existing subsidies. "Twice Subsidy – No Subsidy" compares a regime with twice the baseline expected subsidy intensity to the no-subsidy counterfactual. "High Accountability – Baseline" compares the high-accountability counterfactual with the baseline political and subsidy environment.

than a small effect of investment on pollution.[41] Doubling expected grant intensity increases the pollution reduction to roughly 1% in the full sample and 3–4% among high-pollution systems. These improvements remain substantially smaller than those generated by the high-accountability counterfactual, which reduces pollution among high-pollution systems by approximately 6–9% of average pollution. Expanded subsidies improve water quality by lowering the share of investment costs passed through to residents, but weak accountability limits how strongly these gains translate into additional investment.

8.2.2. *Mandating investment.* To contrast with the subsidy counterfactuals, which preserve local discretion, we implement a stylized mandate requiring every system to

[41]Recall that, conditional on investment, pollution falls by approximately 30% three to eight years later, as shown in Figure 1(A). By contrast, relative to the no-subsidy counterfactual, the baseline subsidy regime increases the investment probability among systems in the highest quintile of underlying pollution by approximately 0.8 percentage points.

Figure 6. Pollution and Welfare Effects of Investment Mandates

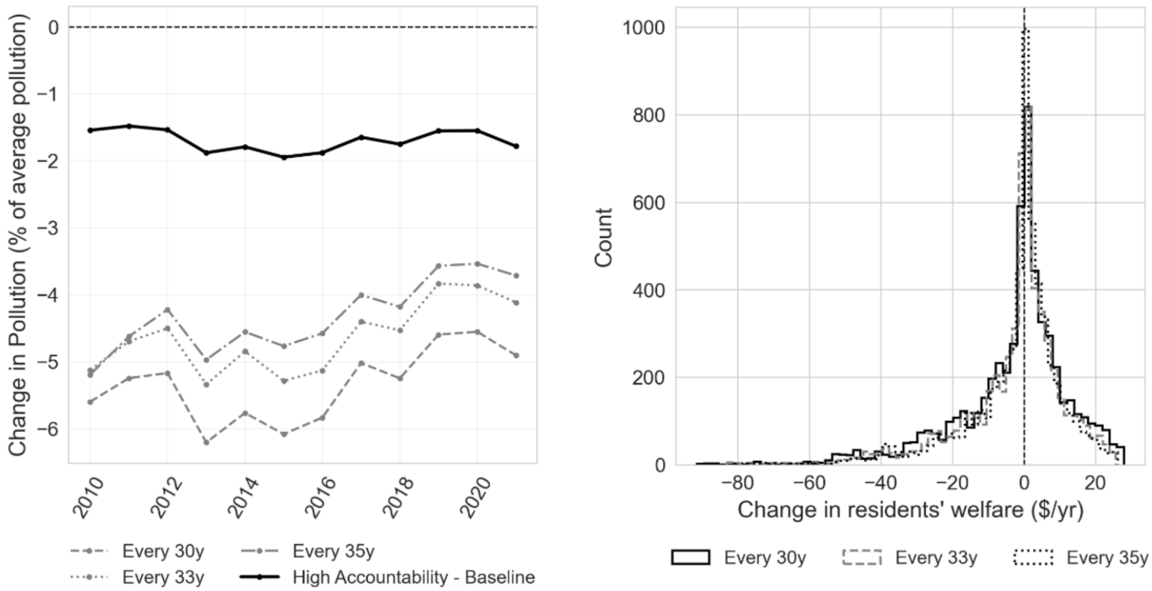

(A) Changes in Aggregate Pollution

(B) Changes in Resident Welfare

Notes: Panel (A) reports the annual change in aggregate pollution under counterfactuals requiring every system to replace its infrastructure every 30, 33, or 35 years, relative to the baseline, expressed as a percentage of the cross-jurisdiction standard deviation of pollution and averaged across simulated paths. The effect of the high-accountability counterfactual relative to the baseline is shown for comparison. Panel (B) reports the distribution of changes in residents' annual welfare, in 2020 dollars, under each mandate relative to the baseline. Negative values indicate that the increase in water rates exceeds the resident's valuation of the associated improvement in water quality.

make investments on its infrastructure on a fixed schedule. We consider investment schedules of every 30, 33, and 35 years, calibrated so that the resulting increase in investment among high-pollution systems is comparable to the increase those systems undertake voluntarily under the high-accountability counterfactual.[42]

Figure 6(A) shows that each mandate produces a larger reduction in aggregate pollution than the high-accountability counterfactual. The reason is mechanical: greater accountability selectively increases investment where its net welfare returns are high, whereas a mandate requires investment across all systems, irrespective of its costs. A shorter replacement cycle therefore produces larger aggregate pollution reductions by inducing more investment, including in jurisdictions where greater accountability alone would not lead officials to invest.

[42]Before recording outcomes for 2003–2022, we simulate the model for 10 additional years, allowing the distribution of investment to converge to its stationary distribution.

This broader pollution reduction comes at a welfare cost because a common replacement schedule does not account for heterogeneity in baseline pollution, financing costs, or residents' willingness to pay for cleaner water. As a result, some systems are required to invest even when the resulting increase in water rates exceeds residents' valuation of the improvement in water quality. Panel 6(B) displays the distribution of annual resident welfare changes. Although the mandates generate substantial gains in some systems, approximately one-third of residents experience a welfare loss under each replacement schedule.

9. Conclusion

We build and estimate a model to quantify the distortions in infrastructure provision that arise when local elected officials are insufficiently responsive to the benefits investment generates for their constituents. Applying our framework to drinking water in California, we find that limited electoral accountability substantially impairs the targeting of investments, with the largest shortfalls concentrated in high-pollution and disadvantaged communities. Conventional policy remedies imperfectly offset these distortions: subsidy take-up remains poorly aligned with investment returns, while uniform investment mandates can generate large welfare losses. Beyond drinking water, our results suggest that, while decentralization is often justified by local governments' superior information about heterogeneous community needs (Oates, 1972), this informational advantage translates into efficient targeting only when policymakers are electorally disciplined. Thus, an important direction for future research is to examine how infrastructure funding and oversight policies should be designed when local political institutions remain imperfect, as well as whether electoral reforms can strengthen the link between resident welfare and infrastructure investment.

References

Agrawal, Ashwini K and Daniel Kim, "Municipal bond insurance and public infrastructure: Evidence from drinking water," *Available at SSRN 3813348*, 2022.

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.

Almagro, Milena and Tomás Domínguez-Iino, "Location sorting and endogenous amenities: Evidence from amsterdam," *Econometrica*, 2025, *93* (3), 1031–1071.

Anderson, Michael L, “As the wind blows: The effects of long-term exposure to air pollution on mortality,” *Journal of the European Economic Association*, 2020, *18* (4), 1886–1927.

Andrews, Donald W. K., “Tests for Parameter Instability and Structural Change with Unknown Change Point,” *Econometrica*, 1993, *61* (4), 821–856.

Ariyarathne, Gangani and Alexander C. Nwala, “3DLNews: A Three-decade Dataset of US Local News Articles,” in “Proceedings of the 33rd ACM International Conference on Information and Knowledge Management” ACM New York, NY, USA 2024, pp. 1–5.

Avis, Eric, Claudio Ferraz, and Frederico Finan, “Do government audits reduce corruption? Estimating the impacts of exposing corrupt politicians,” *Journal of Political Economy*, 2018, *126* (5), 1912–1964.

Balazs, Carolina, Rachel Morello-Frosch, Alan Hubbard, and Isha Ray, “Social disparities in nitrate-contaminated drinking water in California’s San Joaquin Valley,” *Environmental Health Perspectives*, 2011, *119* (9), 1272.

Barwick, Panle Jia, Shanjun Li, Andrew Waxman, Jing Wu, and Tianli Xia, “Efficiency and equity impacts of urban transportation policies with equilibrium sorting,” *American Economic Review*, 2024, *114* (10), 3161–3205.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan, “A unified framework for measuring preferences for schools and neighborhoods,” *Journal of political economy*, 2007, *115* (4), 588–638.

———, Robert McMillan, Alvin Murphy, and Christopher Timmins, “A dynamic model of demand for houses and neighborhoods,” *Econometrica*, 2016, *84* (3), 893–942.

Berry, S, J Levinsohn, and A Pakes, “Automobile prices in market equilibrium,” *Econometrica*, 1995, *63* (4), 841–890.

Besley, Timothy and Anne Case, “Does electoral accountability affect economic policy choices? Evidence from gubernatorial term limits,” *The Quarterly Journal of Economics*, 1995, *110* (3), 769–798.

Biasi, Barbara, Julien Lafortune, and David Schönholzer, “What works and for whom? Effectiveness and efficiency of school capital investments across the US,” *The Quarterly Journal of Economics*, 2025, *140* (3), 2329–2379.

Bishop, Kelly C and Alvin D Murphy, “Valuing time-varying attributes using the hedonic model: when is a dynamic approach necessary?,” *Review of Economics and Statistics*, 2019, *101* (1), 134–145.

Bordeu, Olivia, "Commuting infrastructure in fragmented cities," 2023.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou, "Investment in infrastructure and trade: The case of ports," 2024.

Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel, "The value of democracy: evidence from road building in Kenya," *American Economic Review*, 2015, *105* (6), 1817–1851.

Cellini, Stephanie Riegg, Fernando Ferreira, and Jesse Rothstein, "The value of school facility investments: Evidence from a dynamic regression discontinuity design," *The Quarterly Journal of Economics*, 2010, *125* (1), 215–261.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The effect of minimum wages on low-wage jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

de Benedictis-Kessner, Justin, "Off-cycle and out of office: Election timing and the incumbency advantage," *The Journal of Politics*, 2018, *80* (1), 119–132.

Dobbin, Kristin, Amanda Fencl, and Justin McBride, "California Community Water System Institutional Type Dataset," `https://doi.org/10.25338/B8KP92` 2023.

Fajgelbaum, Pablo D, Cecile Gaubert, Nicole Gorton, Eduardo Morales, and Edouard Schaal, "Political Preferences and the Spatial Distribution of Infrastructure: Evidence from California's High-Speed Rail," 2023.

Gowrisankaran, Gautam, Ashley Langer, and Mar Reguant, "Energy transitions in regulated markets," *American Economic Review*, 2026, *116* (8), 2928–2961.

Hajnal, Zoltan L and Paul G Lewis, "Municipal institutions and voter turnout in local elections," *Urban Affairs Review*, 2003, *38* (5), 645–668.

———, Vladimir Kogan, and G Agustin Markarian, "Who votes: City election timing and voter composition," *American Political Science Review*, 2022, *116* (1), 374–383.

Hansen, Katy and Megan Mullin, "Barriers to water infrastructure investment: Findings from a survey of US local elected officials," *PLoS Water*, 2022, *1* (8).

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

Hsiao, Allan, "Sea Level Rise and Urban Adaptation in Jakarta," 2025.

Iaryczower, Matias, Gabriel Lopez-Moctezuma, and Adam Meirowitz, "Career concerns and the dynamics of electoral accountability," *American Journal of*

Political Science, 2024, *68* (2), 696–713.

Kashner, Zane, "Building with Externalities: Local Governments and Wind Farms," 2026.

Keiser, David A, Bhashkar Mazumder, David Molitor, and Joseph S Shapiro, "Water Works: Causes and Consequences of Safe Drinking Water in America," 2023.

Lim, Claire SH and Ali Yurukoglu, "Dynamic natural monopoly regulation: Time inconsistency, moral hazard, and political environments," *Journal of Political Economy*, 2018, *126* (1), 263–312.

Liscow, Zachary, Cailin Slattery, and Will Nober, "State Capacity and Infrastructure Costs," *Yale Law & Econ. Research Paper*, 2025.

List, John A and Daniel M Sturm, "How elections matter: Theory and evidence from environmental policy," *The Quarterly Journal of Economics*, 2006, *121* (4), 1249–1281.

Mian, Atif, Amir Sufi, and Francesco Trebbi, "The political economy of the US mortgage default crisis," *American Economic Review*, 2010, *100* (5), 1967–1998.

Oates, Wallace E., *Fiscal Federalism*, New York: Harcourt Brace Jovanovich, 1972.

Ouazad, Amine and Matthew E Kahn, "Time Consistent Infrastructure Investments: Optimal Flood Protection Policies in Spatial Equilibrium," 2025.

Posenau, Kelly, "Debt and water: Effects of bondholder protections on public goods," 2021.

Rafey, Will, "Droughts, Deluges, and (River) Diversions: Valuing Market-Based Water Reallocation," *American Economic Review*, 2023, *113* (2), 430–471.

Rust, John, "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, 1987, pp. 999–1033.

Sieg, Holger and Chamna Yoon, "Estimating dynamic games of electoral competition to evaluate term limits in us gubernatorial elections," *American Economic Review*, 2017, *107* (7), 1824–1857.

_______ and _______, "Electoral Accountability and Control in US Cities," *Journal of Political Economy*, 2022, *130* (11), 2985–3023.

Sileo, Gretchen, "Proactive and reactive infrastructure investment," *Available at SSRN 4274528*, 2023.

Timmins, Christopher, "Measuring the dynamic efficiency costs of regulators' preferences: Municipal water utilities in the arid west," *Econometrica*, 2002, *70* (2), 603–629.

Supplemental Appendix for

Inefficiencies in Local Infrastructure: Evidence from Drinking Water in California

Appendix A. Drinking Water Quality and Elections

This appendix provides additional evidence on the electoral mechanisms underlying our analysis. We first show that, when there is a local election, poorer drinking water quality increases the likelihood that a challenger wins the election, consistent with voters holding incumbents accountable for water quality. We then provide dynamic evidence around the implementation of the 2015 California Voter Participation Act and examine how moving elections on-cycle affects electoral contestation and incumbent vote margins.

A.1. **Does Drinking Water Quality Affect Election Outcomes?** Our interpretation of electoral accountability requires that voters respond to drinking water quality. We examine this implication by testing whether greater pollution increases the likelihood that an incumbent is defeated. Because an incumbent may not run if winning chances are low, the dependent variable equals one if at least one challenger wins a local election in the jurisdiction.[1]

To examine whether voters respond to drinking water quality, we estimate

$$\text{ChallengerWin}_{j\tau} = \beta\,\text{Pollution}_{j\tau} + \mathbf{X}'_{j\tau}\gamma + \mu_{c(j)} + \lambda_\tau + \varepsilon_{j\tau},$$

where j indexes jurisdictions and τ indexes two-year election periods. The dependent variable, $\text{ChallengerWin}_{j\tau}$, equals one if at least one challenger wins a local election in jurisdiction j during period τ. $\text{Pollution}_{j\tau}$ is either the continuous aggregate pollution measure used in the model estimation (as defined in Section 6.1) or an indicator for whether this measure exceeds the median among observations with positive pollution. The vector $\mathbf{X}_{j\tau}$ contains the demographic controls (i.e., median household income, the shares of White and Hispanic residents, and the shares of residents below age 16 and above age 65), while $\mu_{c(j)}$ and λ_τ denote county and two-year-period fixed effects, respectively.

[1]Candidate incumbency status is available only for elections recorded in the California Elections Data Archive (CEDA), which primarily covers city and county elections. We therefore exclude many special-district elections and other local contests for which we cannot reliably distinguish incumbents from challengers.

Table A1. Voters Respond to Drinking Water Quality

	Any challenger wins an election			
	OLS	2SLS	OLS	2SLS
	(1)	(2)	(3)	(4)
Pollution measure	0.007***	0.027***		
	(0.0027)	(0.0098)		
High pollution reading			0.0211	0.210***
			(0.0220)	(0.0771)
Cragg-Donald Wald F stat.	-	330.661	-	250.910
Kleibergen-Paap rk Wald F stat.	-	48.637	-	70.863
Demographics	Yes	Yes	Yes	Yes
Year FE, County FE	Yes	Yes	Yes	Yes
Observations	3,187	2,943	3,187	2,943
Mean of the dependent variable	0.751	0.757	0.751	0.757

Notes. This table presents both OLS and 2SLS results. The dependent variable equals one if at least one challenger wins a local election in a jurisdiction during a given two-year period. The sample is restricted to elections recorded in the California Elections Data Archive (CEDA), primarily city and county elections, for which candidates can be classified as incumbents or challengers. The continuous pollution measure is the aggregate measure used in the structural estimation. "High pollution reading" equals one when this measure exceeds the median among observations with positive pollution. In the 2SLS specifications, drinking water pollution is instrumented with surface-water pollution. All specifications include year and county fixed effects and control for standardized median household income and the shares of White residents, Hispanic residents, residents below age 16, and residents above age 65. Standard errors are reported in parentheses. The Cragg–Donald and Kleibergen–Paap first-stage statistics are reported below the coefficient estimates. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

A concern with OLS is that observed pollution may be correlated with unobserved determinants of electoral competition. We therefore instrument drinking water pollution with the surface-water pollution associated with a jurisdiction, as constructed in Appendix OA1.3. The identifying assumption is that, conditional on the controls and fixed effects, surface-water pollution affects local election outcomes through its effect on drinking water quality rather than through other channels.[2] The first stages are strong: the Kleibergen–Paap Wald statistics are 48.6 for the continuous pollution measure and 70.9 for the high-pollution indicator.

Table A1 shows that poorer drinking water quality increases the likelihood that a challenger wins. For the continuous pollution measure, the OLS estimate implies a 0.7-percentage-point increase in the probability of a challenger victory, while the

[2]For example, higher contaminant concentrations in untreated source water require more intensive treatment (e.g. Price and Heberling, 2018, 2020).

2SLS estimate is larger at 2.7 percentage points; both are statistically significant at the one-percent level. Using the high-pollution indicator, the OLS estimate is small and imprecise, whereas the 2SLS estimate indicates that a high-pollution reading raises the probability that at least one challenger wins by approximately 21 percentage points, relative to a sample mean of about 76 percent. Overall, the results suggest that severe drinking water problems materially increase the electoral risk faced by incumbents.

A.2. **Dynamic Event Study: The 2015 CA Voter Participation Act.** To examine the timing of the estimated effect and assess the identifying parallel-trends assumption, we estimate a dynamic difference-in-differences specification centered on the enactment of SB 415.

A.2.1. *Identify changes in election timing.* First, we need to identify changes in election timing. To this end, we construct a biennial indicator of whether each local government follows an on-cycle election schedule. We classify a two-year period as on-cycle only when all elections observed during that period are on-cycle; if any off-cycle election occurs, we classify the period as off-cycle. We retain one observation for each two-year election cycle, indexed by the even calendar year, and exclude jurisdictions for which the election records are too sparse to identify changes in the election schedule reliably.[3]

For each jurisdiction, we test for a single discrete change in the on-cycle indicator at an unknown two-year period. For every candidate break period τ within the central 70 percent of the jurisdiction's time series, we estimate

$$\text{OnCycle}_{it} = \alpha_i + \delta_i(\tau)\mathbf{1}\{t > \tau\} + \varepsilon_{it}$$

and compute the heteroskedasticity-robust F-statistic for the null hypothesis that $\delta_i(\tau) = 0$. Following the unknown-break-point procedure of Andrews (1993), we define the Quandt Likelihood Ratio statistic as the largest of these F-statistics across the candidate break periods and estimate the break period as the two-year period at which this maximum occurs.[4] We classify a jurisdiction as switching from off-cycle to on-cycle elections when the maximum statistic exceeds the 5 percent critical value

[3]When no election is observed for a jurisdiction in a given two-year period,we cannot determine whether the jurisdiction was scheduled to hold an on-cycle or off-cycle election. We therefore restrict the analysis to jurisdictions for which we observe elections in at least four two-year periods during the study period.

[4]When more than one candidate period produces the same maximum statistic, we select the latest two-year period.

FIGURE A1. Dynamic Effects of Exposure to SB 415

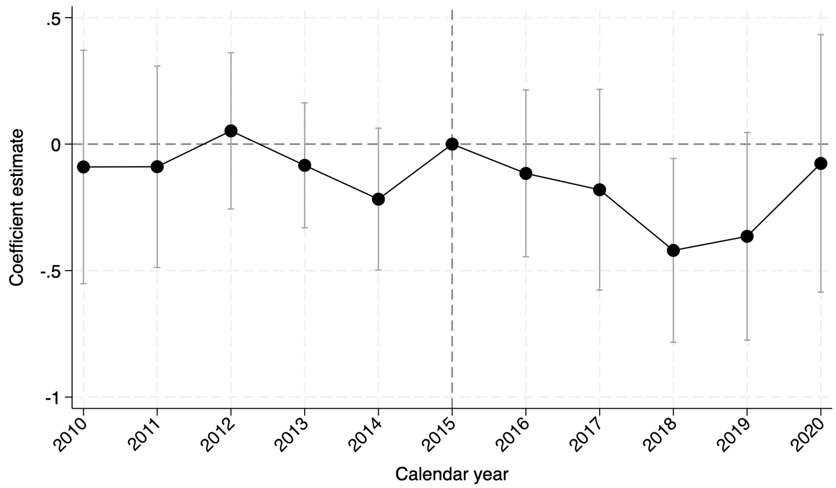

Notes: This figure reports coefficients from a dynamic difference-in-differences specification comparing jurisdictions that switched from off-cycle to on-cycle elections in the 2016, 2018, or 2020 election cycles with jurisdictions whose election-timing regimes did not change after the enactment of the 2015 California Voter Participation Act (SB 415). The omitted reference year is 2015. The regression is estimated by Poisson pseudo-maximum likelihood and includes pollutant-by-jurisdiction and pollutant-by-year fixed effects and the controls listed in Table 2. Vertical bars denote 95 percent confidence intervals based on standard errors clustered at the local-government level.

and the estimated change is positive. Jurisdictions without a statistically significant positive break are classified as following a constant regime: off-cycle if their average on-cycle share is below one-half and on-cycle if it is above one-half.

Given this approach, we identify 59 jurisdictions that remained off-cycle, 24 that switched on-cycle before 2015, and 42 that switched after 2015.

A.2.2. *Event Study.* The treated group for this exercise consists of jurisdictions whose estimated election regime changes from off-cycle to on-cycle after 2015; the comparison group consists of jurisdictions with no post-2015 transition. We interact membership in this group with calendar-year indicators and omit 2015 as the reference year. We estimate the model using Poisson pseudo-maximum likelihood and include the same pollutant-by-jurisdiction fixed effects, pollutant-by-year fixed effects, controls, and jurisdiction-level clustered standard errors as in the main specification.

Figure A1 plots the estimated year-specific coefficients and 95 percent confidence intervals. The estimates for 2010 through 2014 fluctuate around zero and are individually statistically insignificant, providing no evidence of a systematic differential

Table A2. Election Timing and Contested Elections

	Any contested elections			
	Full sample		Subsample	
	(1)	(2)	(3)	(4)
On-cycle election schedule	0.125***	0.117***	0.247***	0.235***
	(0.0326)	(0.0330)	(0.0474)	(0.0506)
Time-varying controls	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes
Jurisdiction FE	Yes	Yes	Yes	Yes
Observations	7,039	6,926	3,230	3,179
R^2	0.562	0.565	0.604	0.606

Notes. This table reports regression results where the dependent variable is an indicator for whether an election is contested. The unit of observation is jurisdiction × two-year period, reflecting that elections typically occur every two years. Standard errors clustered at the jurisdiction level are reported in parentheses. Statistical significance is indicated as follows: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Columns (1)–(2) use the full sample, while Columns (3)–(4) restrict the data to the post-2015 period and excludes jurisdictions that already had on-cycle election schedules prior to 2018, thereby exploiting variation from the implementation of the California Voter Participation Rights Act (SB 415, 2015). Controls include the logarithm of median household income, population, and population density; the shares of White and Hispanic residents and homeowners; and measures of precipitation levels and intensity.

pre-trend among jurisdictions that subsequently switched. The estimates become negative after enactment and grow in magnitude as implementation proceeds. They reach −0.420 in 2018 and −0.365 in 2019. The estimate moves back toward zero in 2020 and is imprecisely estimated.

A.3. **Effects of Election Timing.** Tables A2 and A3 provide additional evidence on how election timing affects electoral competition in our data. Table A2 examines whether elections are contested. Using jurisdiction × two-year observations with jurisdiction and year fixed effects, we find that elections held on-cycle—that is, aligned with statewide or federal contests—are substantially more likely to attract multiple candidates. Across specifications, moving to an on-cycle schedule increases the probability that an election is contested by roughly 12 to 25 percentage points. These patterns are robust to the inclusion of demographic controls and also hold in a subsample focusing on jurisdictions affected by the implementation of the 2015 California Voter Participation Rights Act (SB 415).

Table A3 examines electoral competitiveness using incumbent vote margins. The dependent variable is defined as the maximum of zero and the incumbent vote margins averaged across elections within a jurisdiction $\times$ two-year period. Columns (1) and (2) use the unbalanced sample, which includes only elections for which vote margins are observed. Columns (3) and (4) use a balanced measure that assigns a margin of one to uncontested elections identified in the CEDA database. In the balanced specifications, on-cycle elections are associated with incumbent margins that are approximately 9.6 to 15.2 percentage points smaller, with the estimates statistically significant at conventional levels. By contrast, the estimates in the unbalanced sample are smaller and imprecise. Thus, the relationship between election timing and incumbent margins appears to be driven primarily by the extensive margin of whether an election is contested, rather than by changes in vote margins conditional on a contested race.

This distinction helps reconcile our findings with de Benedictis-Kessner (2018), who finds that on-cycle mayoral elections generate a larger incumbency advantage. Their analysis conditions on elections that are already contested and is therefore most comparable to our unbalanced specifications. In those columns, we find no robust relationship between election timing and incumbent margins; indeed, the positive point estimate in column (1), although close to zero, has the same sign as the result in de Benedictis-Kessner (2018). The negative and statistically significant estimates arise only in the balanced specifications, which incorporate uncontested elections and therefore capture the broader effect of election timing on electoral competition.

Appendix B. Resident Preferences for Water Quality

This appendix provides additional evidence and implementation details for estimating resident preferences for water quality. We first present an event study of house prices following health-based violations and hedonic regressions, which provide evidence that residents value drinking water quality and complement our boundary-design estimates. We then describe the construction of the demand data, computational implementation, and house-price instrument, provide additional evidence supporting the validity of the boundary design, and report further results on heterogeneity in willingness to pay.

B.1. **Complementary evidence.** As additional evidence of residents' distaste for drinking water pollution, we provide estimates from an event study of house-price changes within a water system following a health-based violation. This design uses only within-entity variation and thus corroborates the boundary-design estimates with a different source of identifying variation. Denoting a water system by j with

Table A3. Election Timing and Vote Margins

	Incumbents' vote margins			
	Unbalanced		Balanced	
	Full	Subsample	Full	Subsample
	(1)	(2)	(3)	(4)
On-cycle election schedule	0.00487	-0.0739	-0.0955*	-0.152**
	(0.0295)	(0.0573)	(0.0502)	(0.0687)
Time-varying controls	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Jurisdiction FE	Yes	Yes	Yes	Yes
Observations	2,845	1,155	3,671	1,605
R^2	0.555	0.674	0.474	0.611

Notes. The dependent variable is $\max\{0, \text{incumbent vote margin averaged across elections}\}$. The unit of observation is a jurisdiction $\times$ two-year period, reflecting that elections typically occur every two years. Standard errors clustered at the jurisdiction level are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Columns differ by sample definition. "Full" uses the entire sample. "Subsample" restricts the data to the post-2014 period and excludes jurisdictions that already had on-cycle election schedules prior to 2018, thereby exploiting variation from the implementation of the California Voter Participation Rights Act (SB 415, 2015). "Balanced" imputes an incumbent vote margin of 1 in uncontested elections when election information is available from the CEDA database, while "Unbalanced" does not perform this imputation. Controls are identical to those of Table A2.

violation event e in year t, and a house by h, we estimate

$$\log p_{hjet} = \beta x_{ht} + \phi_{je} + \mu_t + \sum_{\tau} \delta_{\tau}^{I} \ I_{j,t,e}^{\tau} \mathbb{1}\{e \text{ is Tier I}\} \qquad (15)$$
$$+ \sum_{\tau} \delta_{\tau}^{II} \ I_{j,t,e}^{\tau} \mathbb{1}\{e \text{ is Tier II}\} + \epsilon_{hjet}.$$

The treatment dummy $I_{j,t,e}^{\tau}$ equals 1 if the violation e occurred τ years from year t by jurisdiction j, and the specification includes event-specific water system fixed effects and year fixed effects, to account for systems with multiple violations over our sample. We control for a rich set of house characteristics x_{hjet}, as in our boundary discontinuity designs. We estimate the dynamic treatment effects δ_{τ} separately for Tier I and Tier II violations, to allow for different levels pollution to impact house prices differently. Moreover, the information available to residents and potential house buyers might be different as Tier I violations have stringent and immediate reporting requirements -residents are notified within 24 hours. We report dynamic treatment

FIGURE B2. Event Study: Effects of Violations on House Prices

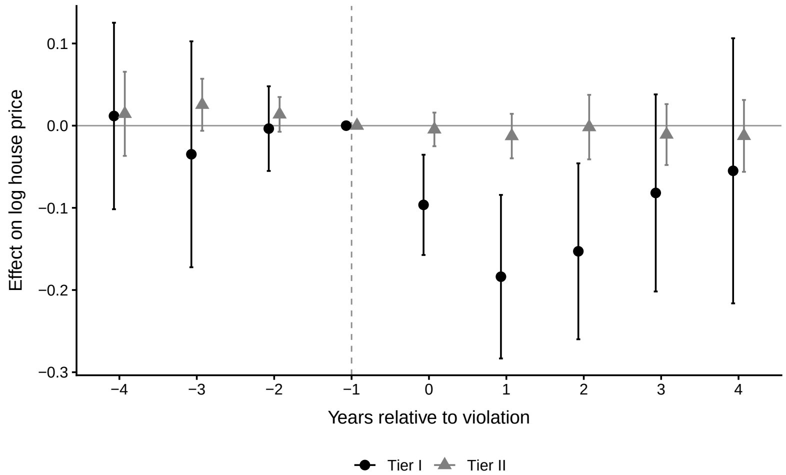

Notes: This figure presents estimates from the event-study specification (15) estimated with a two-way fixed effects estimator. We report point estimates and 95% confidence intervals for Tier I and Tier II violations. Standard errors are clustered at the system-violation event level.

effects estimates in Figure B2. Consistent with our main analysis, we find that Tier I violations adversely impact house prices in the medium run.

Next, Table B4 below reports hedonic estimates of the impact of pollution on house prices, using the same sample of housing transactions that we use to recover residents' preferences in the model. Estimates of these simpler specifications are broadly in line with our preferred estimates, as the annualized effect of a one standard deviation improvement in the pollution index corresponding to column (2) in Table B4 would be \$5,000 $\times$ 0.075 = \$375 per household per year.

B.2. **Implementing the housing demand estimation.** This section provides additional details on the implementation and validation of the housing demand estimation (Section 6.4). We first describe the construction of the housing choice data, the definition of geographic markets and water-system service areas, and the computational approach used to estimate the model. We then describe the construction and relevance of the house-price instrument. Lastly, we provide evidence that house characteristics vary smoothly across water-system boundaries, supporting our boundary design.

B.2.1. *Data construction.* We follow Bayer et al. (2016) and link the CoreLogic housing transaction (deed) dataset to individual mortgage applicant data made available

Table B4. Do Residents Care about Water Quality?

	Housing cost	
	(1) OLS	(2) 2SLS
Drinking water pollution	-5,872.71***	-5,025.24*
	(749.10)	(2,652.60)
Other house and neighborhood char.	Yes	Yes
Boundary × year FE and market FE	Yes	Yes
City × year FE	No	Yes
School × year FE	No	Yes
Drinking-water pollution IV	No	Yes
Observations	1,040,003	969,742
First-stage Wald stat: Pollution		149.84

Notes: Housing cost is the house transaction price plus the net present value (NPV) of water rates over the subsequent 14 years, measured in 2010 dollars. Other house and neighborhood controls include house characteristics (bedrooms, baths, building and lot square footage, year-built bins, land use, and pool) and neighborhood demographics (race/ethnicity shares, college share, and median income). Column (1) reports OLS. Column (2) reports 2SLS using the system-specific surface-water pollution index constructed from nearby ambient surface-water readings. Standard errors, clustered at the boundary-year level, are in parentheses. Statistical significance is indicated by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

by the Home Mortgage Disclosure Act (HMDA). After restricting our attention to residential transactions that can be matched with drinking water systems for which we observe water rates and pollution, and eliminating outliers in transaction prices and housing characteristics, this results in a dataset of 4,230,204 house transactions between years 2003 and 2019 where we also observe buyer income, ethnicity and race, and gender.[5] In addition, we assign each house buyer a simulated workplace using the commuting matrix from the Longitudinal Employer-Household Dynamics (LEHD) data. Given the location of the purchased house, workplaces are drawn so as to match the empirical distribution of tract-to-tract commuting flows in California in 2010. Finally, we group transactions into distinct geographic housing markets defined by core-based statistical areas (CBSAs); Figure B3 shows the resulting markets. Within these markets, houses are associated with their corresponding water system. Importantly, water system boundaries often differ from city boundaries, as illustrated in Figure B4.

[5]Specifically, we drop transactions with prices below the 1 and above the 99 percentiles, and observations with bedrooms, bathrooms, building square footage, and lot square footage above the 99.5 percentile.

FIGURE B3. Geographic Markets used in Demand

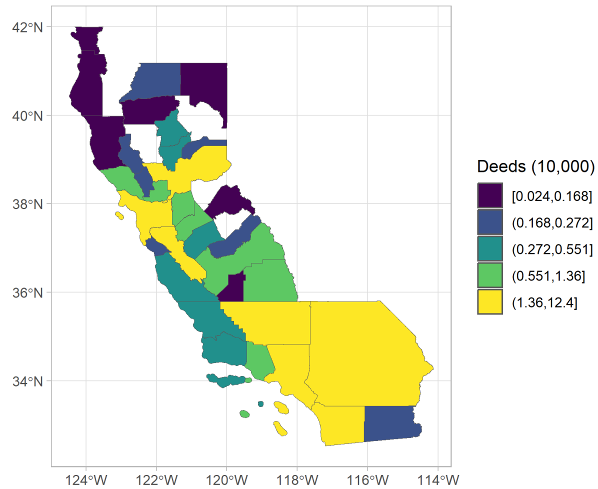

Notes: This figure shows the geographic markets used in housing demand estimation.

FIGURE B4. Boundary Discontinuity Design

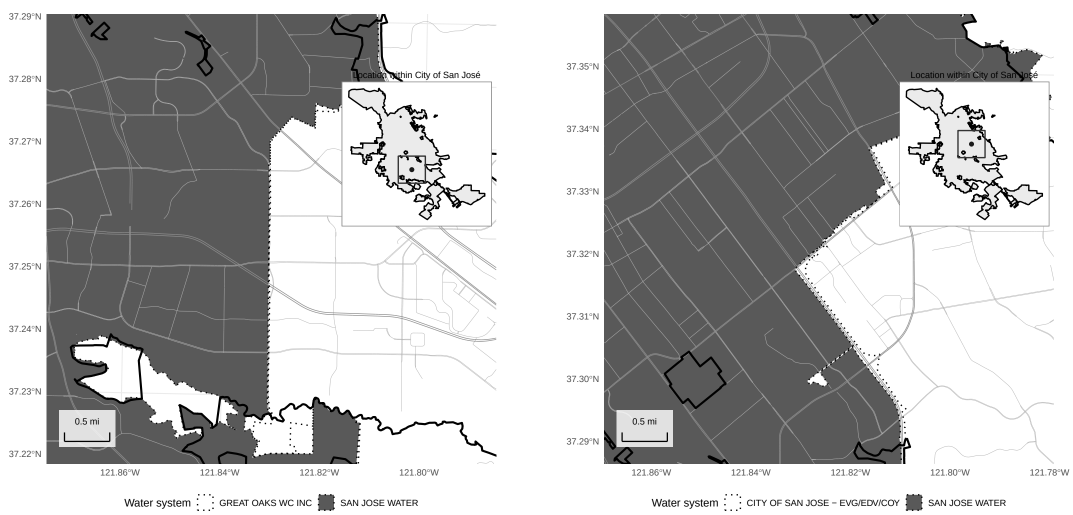

(A) Great Oaks / San Jose Water (B) City of San Jose EEC / San Jose Water

Notes: This figure shows two examples of adjacent water-system service areas whose shared boundaries lie within the City of San Jose. Together, the two panels cover three water systems. The inset in each panel shows the location of the mapped area within the city.

B.2.2. *Computational details.* As shown in Figure B3, the largest number of transactions in the geographic markets considered is above $100,000$. For the largest markets such as Los Angeles and the Bay Area, we define markets as geography–quarter instead of geography–year. This reduces the number of products but the largest market still consists in more than $30,000$ different transactions. As a result, we estimate our demand model using pytorch (Paszke, 2019). Our implementation follows the best practices Conlon and Gortmaker (2020), and we rely on pytorch for GPU acceleration and automatic differentiation. To efficiently apply automatic differentiation of the BLP contraction mapping, we follow the approach used in deep implicit layers models (see for instance, Bai et al. (2019)) and approximate the gradient by using an auxiliary fixed-point mapping that only requires the vector-Jacobian product which is directly provided by automatic differentiation libraries such as pytorch (Paszke, 2019). In doing so, we avoid both inverting a large $J \times J$ matrix and storing in memory all the intermediate tensors used to solve the contraction mapping at every parameter guess. With these changes, we use the L-BFGS optimization routine to estimate the parameters of $(\{\delta_{ht}\}_h, \theta_{qy}, \theta_{py}, \theta_d)$ in Equation (13). We then estimate the demand model described in Section 6 on GPUs.

B.2.3. *House price instrumental variable.* To construct an instrument for house prices, we follow the standard approach in differentiated product demand estimation, as described in Gandhi and Houde (2019), and compute differences in exogenous house characteristics $\boldsymbol{x_h}$–number of bedrooms, bathrooms, building and land square footage, year built. As in Bayer et al. (2007), we exclude nearby houses to form a "donut" of competitors to build price shifters: for a given characteristic ℓ we compute $z^1_{j\ell} = \sum_{k,1\text{km}\leq d_{hk}\leq 5\text{km}}(x_{j\ell} - x_{k\ell})$ and $z^2_{j\ell} = \sum_{k,1\text{km}\leq d_{hk}\leq 5\text{km}}(x_{j\ell} - x_{k\ell})^2$. Finally, we compute the best predictor of house prices, $\hat{p}_{jt} = \mathbb{E}[p_{jt}|\boldsymbol{x}_{jt}, \boldsymbol{z}^1_{jt}, \boldsymbol{z}^2_{jt}]$ using gradient boosted trees and cross-fitting. We show in Figure B5 that this instrument has a relevant first stage, and that information on houses around house j, excluding the ones directly in the vicinity of j, provides predictive power over and above the characteristics of house j itself.

B.2.4. *Continuity of house attributes around the boundary.* Our boundary discontinuity design compares houses located close to one another but served by different public water systems. A potential concern is that water-system boundaries may coincide with discontinuous changes in the housing stock, in which case differences in house prices across the boundary could partly reflect differences in housing quality rather than drinking water quality. Following Bayer et al. (2007), we assess this concern by examining whether observable house characteristics vary discontinuously at system boundaries.

Figure B5. Best Predictor of House Price

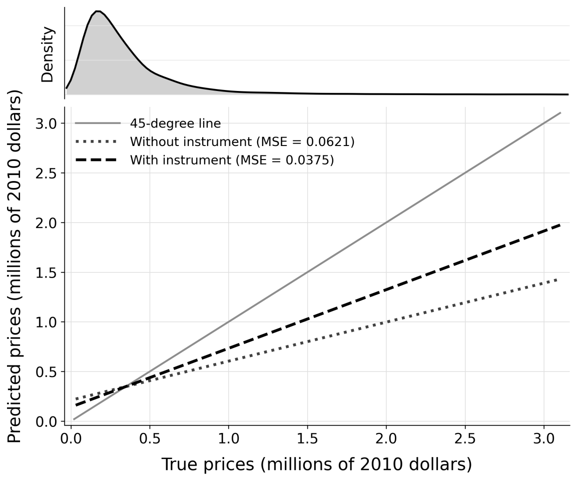

Notes: This figure shows both the line of best fit for the best predictor of house prices based on own-house characteristics only (dotted), and based on own-house characteristics and the first and second order differentiation instruments (dashed) (Gandhi and Houde, 2019) described in the paragraph above. The addition of these instruments improves the fit. The top-panel displays the density of transaction prices.

Figure B6 plots four key housing characteristics—the number of bedrooms and bathrooms, land square footage, and building square footage—against distance to the water-system boundary. Across all four outcomes, we find little evidence of a discrete change at the boundary. The similarity of houses immediately across system boundaries supports the identifying premise of our design that nearby homes served by different water systems are otherwise comparable, and complements the rich set of house characteristics and local fixed effects included in our demand specification.

B.3. **Heterogeneous WTP for Water Quality.** We use the heterogeneous preference estimates in Table 3 to recover each household's marginal willingness to pay for drinking water quality. Under our preferred specification in Column (3) of Table 3, we find that the pollution and housing-cost sensitivities are negative over the relevant range of income, implying that the willingness to pay for water quality is positive.

Figure B7 shows the resulting heterogeneity. The average WTP for a one-standard-deviation reduction in pollution is approximately $356 per household per year. We find that WTP increases with household income. This pattern reflects the two offsetting income interactions in Table 3: higher-income households are less sensitive both to drinking-water pollution and to housing costs. The decline in sensitivity to housing costs is quantitatively stronger, however, so the ratio of the marginal

Figure B6. House Attributes are Continuous across Boundaries

(a) Bedrooms

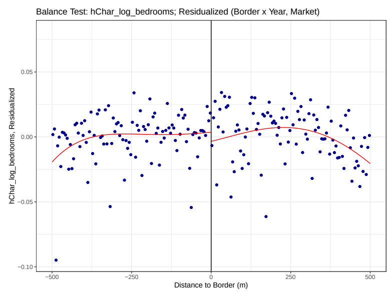

(b) Bathrooms

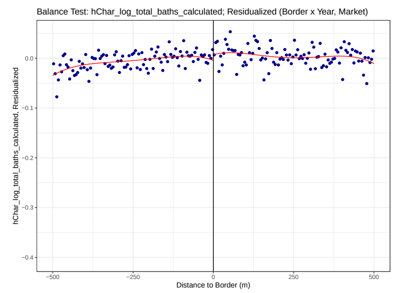

(c) Land Square Footage

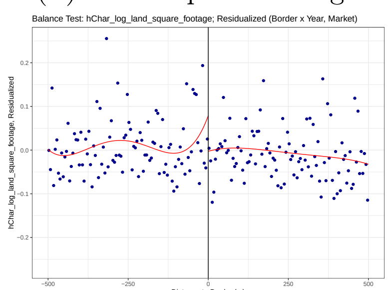

(d) Building Square Footage

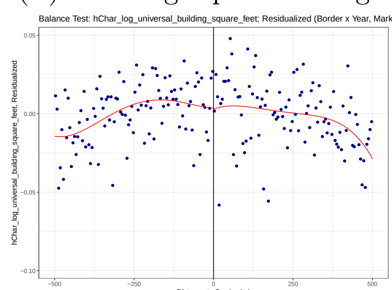

Notes: This figure reports regression discontinuity plots of house characteristics around public water-system boundaries, constructed using the `rdplot` command from the `rdrobust` package (Calonico et al., 2014). The horizontal axis measures distance to the system boundary, with zero denoting the boundary. Distance is normalized so that negative values correspond to the side with higher drinking-water pollution and positive values to the side with lower pollution. The four panels report the number of bedrooms, number of bathrooms, land square footage, and building square footage, respectively.

utility of water quality to the marginal utility of money—and hence WTP—increases with income.

FIGURE B7. Willingness to Pay for Water Quality and Income

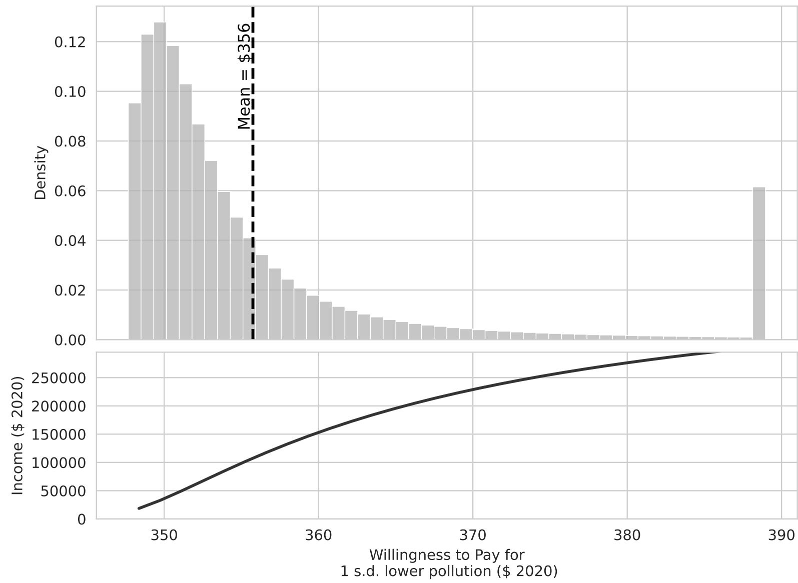

Notes: The top panel plots the distribution across households of annual willingness to pay (WTP), in 2020 dollars, for a one-standard-deviation reduction in drinking-water pollution. For each household, WTP is calculated as the ratio of its estimated marginal utility of pollution to its estimated marginal utility of annualized housing cost, multiplied by one standard deviation of the pollution measure. The calculation uses the preferred estimates in Column (3) of Table 3. The dashed vertical line denotes the sample mean of $356. The bottom panel plots average household income within bins of WTP.

References

Andrews, Donald W. K., "Tests for Parameter Instability and Structural Change with Unknown Change Point," *Econometrica*, 1993, *61* (4), 821–856.

Bai, Shaojie, J Zico Kolter, and Vladlen Koltun, "Deep equilibrium models," *Advances in neural information processing systems*, 2019, *32.*

Bayer, Patrick, Fernando Ferreira, and Robert McMillan, "A unified framework for measuring preferences for schools and neighborhoods," *Journal of political economy*, 2007, *115* (4), 588–638.

———, **Robert McMillan, Alvin Murphy, and Christopher Timmins**, "A dynamic model of demand for houses and neighborhoods," *Econometrica*, 2016, *84* (3), 893–942.

Calonico, Sebastian, Matias D Cattaneo, and Rocio Titiunik, "Robust nonparametric confidence intervals for regression-discontinuity designs," *Econometrica*, 2014, *82* (6), 2295–2326.

Conlon, Christopher and Jeff Gortmaker, "Best practices for differentiated products demand estimation with pyblp," *The RAND Journal of Economics*, 2020, *51* (4), 1108–1161.

de Benedictis-Kessner, Justin, "Off-cycle and out of office: Election timing and the incumbency advantage," *The Journal of Politics*, 2018, *80* (1), 119–132.

Gandhi, Amit and Jean-François Houde, "Measuring substitution patterns in differentiated-products industries," *NBER Working paper*, 2019, (w26375).

Paszke, A, "Pytorch: An imperative style, high-performance deep learning library," *arXiv preprint arXiv:1912.01703*, 2019.

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.

——— **and** ———, "The effects of agricultural and urban land use on drinking water treatment costs: An analysis of United States community water systems," *Water economics and policy*, 2020, *6* (04), 2050008.