Proactive and Reactive Infrastructure Investment*

Gretchen Sileo†
Temple University

August 14, 2023

Abstract

Maintaining infrastructure requires investment. Faced with uncertain quality degradation, managers can invest proactively to prevent failure or reactively to address problems. Using a new dataset on drinking water systems, I estimate a dynamic discrete choice model of infrastructure investment. Simulation results indicate that current investment is insufficient to prevent system deterioration. Proactive-promoting policies facilitate the prevention of most health-based violations but leave some systems vulnerable to extreme quality decline. By contrast, reactive-promoting policies lead to milder but pervasive violations. A policy that increases proactive investment and reserves a safety net of reactive support enables all managers to sustain functional infrastructure.

JEL Codes: C61, H54, Q25

Keywords: Infrastructure Investment, Drinking Water Systems, Dynamic Discrete Choice

*I am grateful to my advisors, Nathan Miller and John Rust, for their invaluable guidance and support. I owe special thanks to Bill Pauley, Don Schierer, Thomas Schubert, Alicia Jacobs, and countless others from the Commonwealth of Kentucky for their assistance in fulfilling my FOIA requests and answering my questions about water quality in Kentucky and the Water Resource Information System. I thank Yunan Ji, Ashley Langer, Will Wheeler, Chris Timmins, Chenyu Yang, Arik Levinson, and seminar participants at multiple institutions for helpful comments. Researcher's own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are those of the researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

†Temple University, Department of Economics, Gladfelter Hall 1115 W. Polett Walk, Philadelphia PA 19122. Email: gretchen.sileo@temple.edu.

1 Introduction

Infrastructure investments fund projects that maintain the quality of existing systems or, when necessary, construct new ones.[1] However, quality can be difficult to observe and measure, and the pace at which systems deteriorate can be uncertain. When allocating an infrastructure budget, decision makers carefully choose between proactive and reactive investment. Proactive investment is made in advance of infrastructure issues. By contrast, reactive investment is made in response to signs of quality deterioration and often becomes necessary when systems are no longer functioning properly. Planned proactive investments prevent predictable failures but at the expense of prematurely replacing components. A reactive investment strategy allows managers to invest only in areas that are no longer operational but at an increased risk of system failure.

In this paper, I examine the proactive and reactive infrastructure investment decisions of community water system managers. Community water systems provide drinking water to over 310 million Americans, or about 94% of the population. In 2020, the Environmental Protection Agency (EPA) found that 7% of community water systems reported at least one health-based violation. Each year roughly 19.5 million cases of waterborne illnesses can be attributed to contaminants in drinking water provided by water systems (Reynolds et al., 2008). The American Society of Civil Engineers (ASCE) has scored America's drinking water infrastructure well below average since 1998 and describes it as "aging and underfunded." The ASCE's Infrastructure Report Card highlights increased costs to maintain and operate systems as well as insufficient funding as significant challenges that water systems face:

> "Maintenance costs reached an all-time high of $50.2 billion above capital in 2017, in part due to deferred capital projects. A recent survey found that 47% of the maintenance work undertaken by utilities is reactive and done as systems fail."

Motivated by these facts, I construct a dynamic discrete choice model of system manager investment in the spirit of Rust (1987) to study managers' proactive and reactive decisions and to determine whether infrastructure investments are sufficient for systems to provide safe drinking water now and in the future. The model is consistent with a number of empirical relationships that emerge from an analysis of data on infrastructure projects and water quality violations in Kentucky over 2007-2019. I incorporate these empirical relationships into the model and quantify the effectiveness of proactive and reactive projects. My findings indicate that proactive projects increase quality more than reactive projects for the same level of expenditure. I also determine that current levels of investment are insufficient to maintain quality above health-based standards. Among a set of policy prescriptions that I consider, policies that only promote proactive investment leave some systems vulnerable to irreversible emergencies, while policies that only promote reactive investment lead to less severe but more widespread violations. A policy that combines modest proactive

[1] Infrastructure is expensive. In 2017 alone the United States public spent $441 billion on transportation and water infrastructure (Congressional Budget Office, 2018).

project increases with substantial reactive project increases provides incentives for managers to invest primarily in effective proactive projects, knowing that there is a safety net of reactive funds in the event of an unexpected disaster, and enables the long term maintenance of safe levels of infrastructure quality.

The paper proceeds as follows. Section 2 provides background information on United States drinking water standards and introduces the data sources. I collect a novel dataset on community water systems and infrastructure projects in the Commonwealth of Kentucky and use these data for my analyses. Kentucky provides an advantageous setting due to the existence of the Water Resource Information System (WRIS), which serves as a comprehensive registry for water systems and infrastructure projects.[2] I use scraping methods to collect over 350 water system reports and more than 3,000 project reports corresponding to the investments pursued by these systems. The reports contain detailed information about the timing, costs, and areas affected by each infrastructure project. I supplement these data with details on violations issued to systems when their water does not meet federal health-based standards. With the ASCE's assessment in mind, I use natural language processing techniques to categorize projects as proactive or reactive.

In Section 3, I develop the four main empirical relationships that motivate the dynamic discrete choice problem. First, managers are more likely to pursue reactive projects if the system experienced a health-based violation in the prior period, implying that managers take action to correct low quality infrastructure issues when they occur. Second, consumers are sensitive to receiving water that violates health-based standards. Using Nielsen scanner data on bottled water sales, I confirm that consumers seek alternative safe drinking water sources when faced with contaminated water at home. Third, the probability of a health-based violation decreases with more proactive project spending, signaling that proactive investments successfully prevent future violations. Lastly, reactive spending successfully decreases the probability of a health-based violation for those systems that spent more time in violation during the prior period. I interpret this as an indication that reactive project expenditures are closely tied to the quality of the system's infrastructure.

Section 4 presents the model developed from these findings. Managers are responsible for maintaining their system's infrastructure quality.[3] Although current quality is known to managers, they do not know exactly how quality will progress in the future. Every period, managers decide to invest or delay investment anticipating the uncertain path that infrastructure quality will take. Investment decisions balance managers' desire to avoid providing consumers with water that violates health-based standards against the cost of quality-improving investments. Without investment, quality reaches progressively lower levels. As quality falls, systems face an increased probability of receiving a health-based violation and managers are more likely to have to undertake a more expensive reactive, as opposed to proactive, project.

[2]Kentucky is advantageous but not entirely unique in terms of publicly available data. Other states, e.g., Wisconsin, host similar drinking water system investment information online.

[3]Infrastructure quality is a measure of the system's structural ability to provide safe drinking water to consumers. Water quality is a measure of the concentration of contaminants present in water provided by the system. See Section 2.1 for a discussion of this distinction.

I discuss the estimation, identification, and results in Section 5. I find the full solution to the model using likelihood methods to determine the parameters that maximize the probability that observed outcomes in the data are predicted by the model. Data on the timing and type of project investments as well as the amount of time systems spend in violation help to pin down the model parameters. The main challenge for estimation is that system quality is both persistent and unobserved by the econometrician. As a result, the standard assumption that unobserved state variables are serially independent cannot be applied in this application, which increases the computational burden of estimation significantly. To make the problem tractable, I apply methods developed in Reich (2018). In particular, I use a combination of fixed-point methods to solve for conditional choice probabilities and recursion methods over the possible quality histories to estimate the model parameters.

The model fits the data well. In simulations, it matches trends for the number of systems that spend time in violation as well as the number of systems undertaking proactive and reactive projects each year. Parameter estimates indicate that reactive projects occur at higher levels of infrastructure quality than violations, implying that managers have the ability to detect potential failures before they occur. Results further reveal that proactive projects are cheaper and reduce quality by a larger amount per investment than reactive projects. Investments are most effective at reducing the amount of time spent in violation when made just prior to violating thresholds. Estimates for the reactive threshold and the optimal timing for investment coincide, indicating that the most opportune timing for managers to undertake projects is right before infrastructure quality crosses the threshold for a reactive project. Managers are unable to maintain quality at these levels due primarily to the low effectiveness of projects at their current sizes. On average, both project types only reduce the amount of time spent in violation by approximately one quarter. This effect is too small to warrant constant investment and is insufficient to offset systems' quality decline.

Section 6 explores counterfactual policies under existing conditions. First, I assess the efficacy of possible penalty and subsidy policy interventions designed to increase the number of projects. I find that these policies can only decrease the average expected amount of time each system spends in violation over 2007-2057 by roughly four years. These simulations demonstrate that relying solely on policies that induce more frequent investment yields limited results because the quality increase from investment is still too low. I next explore counterfactual policies that increase the size of investments. Under these conditions, managers invest in fewer projects and systems spend more time in violation. The results indicate that policies that increase the size or number of projects are individually unable to combat infrastructure quality decline. In subsequent simulations, I combine increased project size with project subsides to analyze outcomes under conditions where managers have a relaxed budget constraint and elevated investment incentives.

I first simulate scenarios for policies that promote either proactive or reactive projects. I find that these policies reduce the average expected time in violation but have different distributional outcomes. Proactive-incentivizing policies enable many managers to avoid violations but expose some systems to long term states of violation. Counterfactual policies that fully subsidize larger reactive projects lead to more systems incurring occasional health-based violations but more equitably spread violation risk across all systems. In a final counterfactual, I simulate a policy that combines proactive and reactive incentives. Results indicate that a policy pairing modest increases in proactive project size with significant reactive project size increases creates conditions in which only 8% of systems violate health-based standards over the next five decades. The few systems that do experience violations spend an average of less than three days per year out of compliance.

I conclude in Section 7 by tying the results of my analysis to current policy discussions and providing suggestions for future research. In 2021, Congress passed the Infrastructure Investment and Jobs Act, which allocates $550 billion for future infrastructure projects and includes support for drinking water systems. My analysis indicates that proactive projects are more efficient than reactive projects but that reactive projects are often undertaken by systems when infrastructure quality is low. Understanding when managers invest and the difference between types of investments can enable the efficient and effective distribution of the funds allocated in the Act. According to my results, a carefully considered incentive program for managers that combines small proactive incentives with larger reactive incentives can lead to higher long term quality levels and reduced violations for all systems as compared to a future without additional federal support.

1.1 Literature Review

I contribute to a literature examining system manager decision making in the context of drinking water. Previous studies, such as Timmins (2002) and Sears et al. (2022), analyze the decisions of managers that have underground aquifers to price water below marginal cost. Both use dynamic structural models of decision making but focus on pricing decisions rather than system managers' long term infrastructure investments. Agrawal and Kim (2022) and Posenau (2022) demonstrate that drinking water system investment is sensitive to financial constraints. Relative to those articles, I prioritize modeling the investment decision process in order to examine outcomes under alternate incentive policies. Keiser et al. (2023) perform a careful analysis of the Safe Drinking Water Act and find evidence that the benefits of drinking water investments outweigh the costs. I use a different approach, taking the importance of drinking water investment as given, that focuses on the dynamic decision making process of managers to understand how they maintain their systems through investment. In my model, managers are averse to incurring an EPA violation, which is supported by the findings of other studies. Grooms (2016) and Bennear and Olmstead (2008) find that public disclosure leads to reductions in violations and Bennear et al. (2009) demonstrate that some managers even strategically alter their behavior to prevent triggering a violation.

Several articles further explore the strategic relationship between regulators and firms that release contaminants into the environment (Abito, 2019; Blundell et al., 2020; Lim and Yurukoglu, 2018; Kang and Silveira, 2021; Leisten and Vreugdenhil, 2023). In these settings, regulators impose standards on profit-maximizing firms that have distortionary effects on firms' incentives. My analysis focuses on publicly owned water systems that aim to provide safe drinking water to consumers at minimal cost. As a result, I am able to capture the economic conditions in my setting with a simpler state space. I can then adapt full solution methods to estimate a model with a continuous serially correlated unobserved state variable, which would be more challenging in a strategic games model.[4] Accounting for unobserved heterogeneity is important for correctly characterizing the economic determinants of infrastructure investment. Using a continuous measure of infrastructure quality as the primary state variable, I am able to estimate smooth functional forms for the consequences of declining quality and ascertain the ability of investment to counteract quality decline. To my knowledge, this paper is one of the first to estimate a model with entirely unobserved states.

Lastly, I add to an operations research literature that studies the optimal timing for repair or replacement of deteriorating systems. Jardine et al. (2006) provides an extensive review of the "condition-based" maintenance literature, which compares strategies of scheduled system component replacement with strategies based on assessments of system condition. Many operations studies assess methods for determining the appropriate timing of infrastructure repairs and discuss best practices for monitoring systems to detect signs of decline (Keizer et al., 2017; Kim and Makis, 2013; Si et al., 2011; Kurt and Kharoufeh, 2010). In my analysis, I assume that managers are aware of their system's infrastructure quality and that they face both per-period uncertainty and uncertainty over the progression of quality. Once per-period uncertainty is realized, I infer investment type from the language system managers use to describe the project. I demonstrate that there are two distinct types of investments, which I call proactive and reactive, and that pointed policy interventions incorporating these distinctions can lead to improved outcomes for consumers.

2 Data and Background

2.1 Drinking Water System Infrastructure

A community drinking water system serves an average of at least 25 customers for the duration of the calendar year. Systems source their water from one of three options: surface water (e.g., lakes or reservoirs), groundwater (e.g., aquifers), or recycled water (e.g., highly treated wastewater). After pumping in water from these sources, systems further filter this water until it is considered safe for human consumption. A water system's infrastructure consists of all the components involved in the movement and treatment of water from the time water enters the system until the point where it reaches the final service connection, e.g., a house or apartment building. The types of structures that typically make up a water system include: pumps to actively move water, pipes for distribution, filtration systems to process water, meters and sensors for measuring contaminants, and tanks for long term storage.

Infrastructure serves as the first line of defense against harmful contaminants. Factors such as the system's age, component materials, and periods of excessive demand can lead to infrastructure

[4] Other articles have integrated unobserved heterogeneity into strategic games using alternative methods that require simulation, instruments, or discretized unobserved states (Ericson and Pakes, 1995; Norets, 2009; Arcidiacono and Miller, 2011; Blevins, 2016; Connault, 2016; Berry and Compiani, 2020; Kalouptsidi et al., 2021).

deterioration and an increased risk of contamination (National Research Council, 2006). Infrastructure maintenance is the responsibility of the system manager and generally entails the replacement or repair of the above structures. Investments that replace nonfunctional pipes or degraded tanks can have a direct impact on the concentration of pollutants present in water provided by the system. Other investments, such as meter replacements or technology upgrades may not alter the concentration of contaminants on their own, but these types of investments enable the manager to better monitor and maintain the quality of the water that the system provides. For the duration of the paper, I examine the decisions of managers to invest in projects to improve infrastructure quality and how these investments relate to the system's ability to provide safe drinking water to consumers.

2.2 United States Water Quality

In 1974, Congress passed the Safe Drinking Water Act to prevent threats to public health by setting standards for acceptable drinking water. The EPA currently sets maximum acceptable levels for 94 different contaminants, which include both chemical and microbial pollutants. System managers rely on contaminant tests and their expertise to determine the proper dosing of treatment additives, the correct application of disinfectant procedures, and the status of filtration absorbent. These types of activities involve the day-to-day operation of the system and are not considered in my analysis. As fluctuations in pollution concentrations have a less direct bearing on infrastructure investment decisions, I summarize the ability of water quality to influence investment with the violation status of a system. A systems' violation status links manager behavior to consumer notification and serves as a motivator for manager investment.

EPA regulations dictate that if a system provides water exceeding the maximum allowable concentration of a contaminant, fails a testing requirement, or provides water posing a health risk, then the system must notify its customers. Customer notifications are classified into three levels based on the severity of the infraction (Environmental Protection Agency, 2009). Tier 1 notifications involve violations that pose the most severe health risks and require notification within 24 hours of discovery. Tier 2 public notification violations are for contaminants that pose a less imminent public health risk and notification is required within 30 days of detection. Tier 3 public notification violations mostly encompass monitoring and reporting violations and customers are often notified through an annual report.[5] Because of the lower severity of these violations, I exclude them from my analysis.

2.3 Data Sources

I collect and aggregate data from multiple sources. Below I provide details on these data sources and the methods I employ for cleaning and compiling the data for use. All prices and costs are

[5]Findings in the Environmental Protection Agency (2008) report "2006 Data reliability analysis of the EPA safe drinking water information system/federal version", indicate that only 30% of these types of violations are reported to the federal Safe Drinking Water Information System.

adjusted to real 2012 dollars using the GDP deflator.

2.3.1 Kentucky Water Systems and Infrastructure Projects

I compile data on Kentucky's water systems and infrastructure projects using information published by the Kentucky Infrastructure Authority on the WRIS online portal. The primary purpose of WRIS is to consolidate information on Kentucky's public water systems for use in water planning, emergency decision making, and to track and allocate funding to support infrastructure projects. There are two main WRIS reports that provide the richest data: the system and project reports. I source additional information regarding projects and the project approval process from historical Intended Use Plans (IUPs) and the Kentucky Infrastructure Authority's Annual Reports, which I obtained through Freedom of Information Act requests to Kentucky's Department of Local Government.

Community Water System Reports

For my analysis, I consider publicly owned community water systems that are active as of 2021.[6] I source information on Kentucky's community water systems from WRIS system reports. These reports contain a county-level breakdown of the population served by the public water system, the date the system was established, the number of employees, the type of water source used, e.g., source water or purchased water, groundwater, or surface water etc., and the exact geographical location of the public water system. More than 95% of all Kentuckians receive water from public water systems. Figures E.1 and E.2 in Appendix E provide the layout of all pipes comprising Kentucky's public water systems and an example of the geographic layout of a public water system.

System reports also contain the median household income for each county population subset served by the system. Relative median household income data is typically calculated using American Community Survey (ACS) data integrated into WRIS. Median household income dictates how Kentucky assigns financial assistance to specific areas through programs such as the Drinking Water State Revolving Fund (DWSRF). Kentucky groups systems and projects into three levels, called Non-Standard Rate Levels (NSRL), which are tied to the Kentucky-wide median household income. These classifications are defined as follows:

- NSRL = 0: Greater than or equal to Kentucky Median Household Income
- NSRL = 1: Between 80% Kentucky Median Household Income and Kentucky Median Household Income (exclusive)
- NSRL = 2: Less than or equal to 80% Kentucky Median Household Income

Table 1 summarizes the WRIS system data. Roughly 82% of systems source their water from surface water sources, which generally require more treatment than groundwater sources due to

[6] I extracted data from WRIS in July 2021. Of the 377 systems with some data at the time of extraction, 361 systems are publicly owned community water systems. This number shrinks to 353 systems when removing those with missing demographic data, e.g., serviceable population, etc.

Table 1: Kentucky Community Water System Statistics

Characteristic	System Count	Mean	St. Deviation	Min	Max
Operational					
Employees	353	12.805	29.215	0	437
Purchaser	185	0.524	0.499	0	1
Surface Water Source	290	0.822	0.383	0	1
Very Small (<500)	12	0.034	0.181	0	1
Small (501 - 3,300)	117	0.331	0.471	0	1
Medium (3,301-10,000)	131	0.371	0.483	0	1
Large (>10,000)	93	0.263	0.441	0	1
NSRL = 0	87	0.247	0.431	0	1
NSRL = 1	102	0.289	0.453	0	1
NSRL = 2	164	0.465	0.499	0	1
Demographic					
Non-white Population	353	0.068	0.048	0.002	0.288
ln(Housing Density)	353	3.509	0.632	2.275	6.728
ln(Med. HH Income)	353	10.575	0.243	10.013	11.395

Notes: Operational statistics are constructed using data drawn from WRIS as of July 2021. Population categories are based on data from the Safe Drinking Water Information System, and the demographic statistics are compiled using the county breakdown percentages from WRIS in conjunction with 2019 county-level data from the ACS.

sediment and pollutant exposure. Most systems provide drinking water to households that fall below Kentucky's median household income (75% of systems have an NSRL categorization of 1 or 2), and approximately 35% of all systems serve populations of less than 3,300.

Infrastructure Project Reports

Investment in water system infrastructure can take many forms, including leaking pipe repairs, installation of newer water meters, and water tank replacements. To initiate an infrastructure project, community water systems must complete a project profile within WRIS outlining the purpose of the project, the expected cost of the project, who is responsible for oversight, the expected timing of the project, and the anticipated funding sources. These project plans are then reviewed and approved by multiple agencies within Kentucky (e.g., the Energy and Environmental Cabinet, Department of Water, Infrastructure Authority, Area Development District Water Management Planning Councils) before they are allowed to proceed. Most project profiles contain complete information and WRIS is used as a "registry" for Kentucky's water infrastructure projects. Information on projects dates to the beginning of WRIS in 2001. Project profiles are also required by Kentucky in order for a project to be eligible to receive financial support from the Drinking Water State Revolving Fund, and many other financial support programs have also adopted this requirement.[7]

In order to perform analysis comparing proactive and reactive projects, I classify each of these reports using natural language processing techniques. I consider a project to be reactive if it contains text indicating that the project is necessary for the system to provide safe drinking water e.g., references to public health emergencies, extensive breaks, overstressing, or unsafe conditions. If a project does not meet these criteria, it is classified as proactive. Details on the classification process and the application of the natural language processing model can be found in Appendix A.

2.3.2 Water System Violations

Water systems are required to test regularly in order to monitor for EPA regulated contaminants and to report failures to meet contaminant standards to the EPA's nationwide database, the Safe Drinking Water Information System (SDWIS). This database includes details on each violation, including the date the violation was first detected, the current status, the triggering event, and the public notification tier. I collect SDWIS violation data over 2006-2019 using the EPA's API.[8]

I restrict my observation period based on the timing of EPA changes to the National Drinking Water Contaminant List and observed water system responses. In 2000, the EPA introduced one new contaminant to the list, and two public water systems experienced compliance issues between 2000-2003. From 2004-2006, there were no water systems in Kentucky with health-based violations. Following the introduction of three new contaminants to the List in 2006, many systems began to violate health-based standards. Approximately 70% of water systems in my data have at least one health-based violation from 2007-2019. I exploit the 2006 change in monitored contaminants to estimate the relationship between infrastructure investment behavior and violations that would otherwise be unidentifiable in a period of complete compliance. In model estimation, I assume that systems are subject to a constant water quality standard for health-based violations over 2007-2019. The EPA did make two revisions to the National Drinking Water Contaminant List during this period, but of the 1,329 violations I observe, only 23 violations—less than 2%—are due to these revisions.

2.3.3 Additional Data Sources

The Nielsen Corporation's Retail Scanner Data (2006-2019), made available to me through the Kilts Center at the University of Chicago, allows me to estimate consumer responses to violations. The scanner dataset contains weekly sales data including pricing and volume for UPCs covering an extensive list of product categories and information from 30,000-50,000 stores across the United States. I reduce this dataset to products with a product module code "Bottled Water," and stores that fall within Kentucky. All observations of powdered additives and filters are removed from the

[7] Information on WRIS obtained from multiple sources within Kentucky, including the "2015 Water Management Plan."

[8] The SDWIS data can be accessed using the EPA's Envirofacts API. The analysis in this paper uses data last accessed in July 2021.

dataset. This leaves observations at 937 stores within Kentucky, covering 112 of the 120 Kentucky counties for at least some part of 2006-2019. There are 352 consolidated brands of bottled water providers in the remaining dataset, including the "control brand," which is a label Nielsen uses to protect an identifiable private label at a specific store (e.g., store-brand bottled water labels). Most frequently, the largest share of bottled water sales in a given week are of the control brand and the average yearly revenue shares for the control brand are 38.9%.

I also employ five-year ACS data covering 2005-2019 to obtain county-level demographics on total population, median household income, and housing units (Manson et al., 2021). I disaggregate this data to the yearly level by averaging values from the surveys covering a given year. Temperature and precipitation data come from the National Oceanic and Atmospheric Administration. I use daily weather station readings from 2006-2019 for stations in and around Kentucky. I construct data for the average maximum temperature, average minimum temperature, and average total precipitation for each county on a weekly basis using the coordinates of the stations to map each county to the closest weather reading.

3 Motivating Results

I present here the findings from my initial descriptive regressions. Results indicate that both system managers and consumers are sensitive to health-based violations: system managers undertake reactive projects and consumers act by increasing bottled water consumption. The behavior of managers indicates that reactive projects are generally undertaken when infrastructure quality is low. I then demonstrate that proactive and reactive infrastructure expenditures successfully reduce the probability of a health-based violation, but act in different ways. In Section 4, I develop a dynamic discrete choice model inspired, in part, by these results.

3.1 System Managers and Consumers Respond to Violations

3.1.1 System Manager Response

First, I explore the timing of projects and whether the probability that a system manager invests in an infrastructure project is linked to the occurrence of a health-based violation. I construct three indicator variables based on (1) whether any type of project is approved in a given year, (2) whether a proactive project is approved, and (3) whether a reactive project is approved. I also create a health-based violation indicator that equals one if a system experienced either a Tier 1 or Tier 2 violation in a given year. Then I run a series of probit regressions to predict the probability of each outcome. The data cover the violation outcomes and infrastructure expenditures of 353 community water systems in Kentucky from 2007-2019.

In this analysis I control for observable system differences including: indicators for the size of the population served by the system, indicators for water source, and the number of employees working at each system. I also use ACS data in combination with the WRIS breakdown of consumers

Table 2: Probability of an Infrastructure Project

	(1)	(2)	(3)
Lag violation	0.063	−0.007	0.195
	(0.067)	(0.075)	(0.077)
Non-white fraction	−0.470	0.015	−0.862
	(0.594)	(0.642)	(0.697)
Small	0.567	0.486	0.481
	(0.172)	(0.191)	(0.229)
Medium	0.746	0.705	0.556
	(0.172)	(0.189)	(0.229)
Large	0.833	0.819	0.628
	(0.174)	(0.191)	(0.231)
Surface Water	0.215	0.178	0.182
	(0.084)	(0.091)	(0.104)
Pseudo R^2	0.068	0.067	0.074
Observations	4,236	4,236	4,236

Notes: The dependent variable is an indicator for any project approved in the year, a proactive project approved in the year, and a reactive project approved in the year for columns (1), (2), and (3) respectively. Additional controls include the natural log of the median household income and housing density, an indicator for whether the system purchases their water, and the number of employees working at the system. The model also includes year and Area Development District fixed effects.

in each county to construct weighted median household income, housing density, and the fraction of the population that is non-white for each year-system observation. I include year fixed effects to control for a variety of possible yearly shocks, including a few natural disasters that damaged water systems during this time period. Lastly, I include a fixed effect for the Area Development District (ADD) in which the system is located. ADDs are geographical groupings within Kentucky that are used for water system coordination and planning.

Table 2 contains the results of these regressions in columns (1), (2), and (3) respectively. As depicted in the table, only reactive projects are more likely to be approved in the year following a health-based violation. I interpret this as evidence that both reactive investments and violations are likely to occur when infrastructure quality is low. System managers are more likely to face increased pressure from their constituencies when providing contaminated water and may invest in a reactive project on the heels of a violation as a response to this pressure.

3.1.2 Consumer Response

One way to determine if consumers care about the violation status of their tap water is to evaluate whether or not consumers seek out alternative safe water sources when exposed to a violation.[9] In

[9] Bottled water facilities are regulated by the FDA and are not subject to the same criteria used to evaluate community water systems. However, when a violation occurs that cannot be fixed by boiling, system managers often either suggest

Table 3: Bottled Water Sales Increase with Health-Based Violations

	Mean	Std. Error
Tier 1 Violation	0.217	0.075
Tier 2 Violation	0.029	0.010
Price per Gallon	−0.490	0.007
Adjusted R^2	0.926	
Observations	77,466	

Data cover 730 weeks from 2006-2019. Of the 120 counties in Kentucky, 112 have sales information for at least one week. Regression includes additional controls for average county precipitation, maximum temperature, and minimum temperature. The model also includes week and county fixed effects.

keeping with other articles that study consumer avoidance behavior, I examine changes in county-level bottled water sales when a system experiences a health-based violation.[10]

I regress the weekly log sales of bottled water at each store (mapped to a county), on the fraction of the week each water system spends in violation. For this analysis, I separately consider the fraction of the year spent in a Tier 1 and Tier 2 violation. The time spent in violation is weighted by the fraction of the county population that is served by the community water system to obtain a more accurate measure of the level of exposure within a county.

The regression equation is given by:

$$y_{ct} = a_0 + a_1 p_{ct} + a_2 \sum_{w \in c} v_{1wt}(n_{wc}/n_c) + a_3 \sum_{w \in c} v_{2wt}(n_{wc}/n_c) + \boldsymbol{b x_{ct}} + d_c + f_t + e_{ct} \tag{1}$$

where variable subscripts c and t refer to the county and the week. The weekly log sales of bottled water is captured by y_{ct}, and the fraction of week water system w spends in a Tier 1 and Tier 2 violation are represented by v_{1wt} and v_{2wt} respectively. The fraction of the county population that is served by water system w is captured by (n_{wc}/n_c). I control for consumer price sensitivity using the county average price per gallon of bottled water, p_{ct}. I use the average maximum and minimum temperature as well as the average total precipitation observed in the county for the week, $\boldsymbol{x_{ct}}$, to control for seasonal changes in demand. I further include a county fixed effect d_c, and a week fixed effect, f_t. The county fixed effects are included to account for differences such as household income or average population. Certain age groups are more susceptible to health effects from exposure to water contaminants, and income levels and age differences might account for variation in averting behavior. The week fixed effects are included to control for time trends that are separate from the climate controls.

or offer bottled water as a substitute. For example, during the public health crisis in Flint, the government sponsored a free bottled water program from January 2016 through April 2018.

[10] Other studies have previously used changes in bottled water sales in different settings to demonstrate consumer sensitivity to water quality (Graff Zivin et al., 2011; Wrenn et al., 2016; Allaire et al., 2019; Christensen et al., 2021).

Table 3 presents the results. Exposure both to a Tier 1 and to a Tier 2 violation results in an increase in bottled water sales, with a 22% and 3% increase respectively. This is consistent with the fact that Tier 1 violations pose the most imminent threat to consumer health and require the quickest turnaround to notify consumers. I also find consumers have a negative response to the cost of a gallon of bottled water indicating that although consumers are sensitive to violations, they are also sensitive to the cost of substitution. From the above analysis, I conclude that consumers within Kentucky are sensitive to changes in tap water quality and substitute to bottled water as an alternative.

3.2 Relationship Between Investment and Violations

Lastly, I use probit regressions to examine the relationship between infrastructure expenditure and the probability of a violation. For this analysis, I construct an additional variable that captures the fraction of the year that a system spent in violation of health-based standards. I use this variable along with the other previously constructed measures to analyze the violation outcomes and infrastructure expenditures of 353 community water systems in Kentucky from 2007-2019.

The novel variable included in the model is the adjusted infrastructure investment calculated from project expenditures approved in the preceding five years.[11] This variable approximates the amount each system has invested in infrastructure improvements. I use this metric because there is often a delay between the timing of a project's approval and the project's construction and installation. I run two regressions to examine the effect of a project on the probability of a health-based violation. In the first regression, I examine the effect of total project investment, regardless of project classification, and in the second regression I examine the separate effects of proactive and reactive project investments. In keeping with the literature, I also control for a series of system attributes known to influence the probability of a violation, including the fraction of the previous year spent in violation.[12]

Table 4 presents the results. The primary variables of interest are the cumulative project expenditures approved in the previous five years, and the interaction between this term and the fraction of the previous year spent in violation. Column (1) depicts the effect of undifferentiated project investment. Systems with more time spent in violation in the prior year and larger project expenditures are less likely to experience a health-based violation in the current period. However, looking at proactive and reactive projects separately, it becomes clear that the two types of projects influence the probability of a health-based violation differently.

Column (2) splits investments into proactive and reactive projects. Wald tests confirm a significant difference between the effects of proactive and reactive projects on their own and when interacted with time spent in violation in the prior period.[13] This supports the initial classification of projects into distinct groupings, as the two types of projects influence the likelihood of a violation

[11] For my analysis, I use the project approval year as the basis for project expenditure adjustment.

[12] Allaire et al. (2018) find evidence that water source, system size, population income, housing density, and other included variables are strong predictors for violations across the United States.

[13] The result of these two tests are $W(1) = 11.85, p < .001$ and $W(1) = 4.55, p < .05$ respectively.

Table 4: Health-Based Violation Probability

	(1)	(2)
Investment ($M in last 5 years)	−0.002	—
	(0.003)	—
Invest. × Lag Violation Fraction	−0.031	—
	(0.016)	—
Proactive Investment ($M in last 5 years)	—	−0.015
	—	(0.005)
Proactive Invest. × Lag Violation Fraction	—	0.028
	—	(0.029)
Reactive Investment ($M in last 5 years)	—	0.008
	—	(0.004)
Reactive Invest. × Lag Violation Fraction	—	−0.064
	—	(0.025)
Lag Violation Fraction	2.419	2.403
	(0.149)	(0.151)
Pseudo R^2	0.242	0.245
Observations	4,236	4,236

Notes: The dependent variable is an indicator for a health-based violation. Additional controls include indicators for the size of the population served by the system, i.e., small, medium, large, as well as the natural log of the median household income and housing density, an indicator for whether the system purchases their water, an indicator for whether the system uses surface water, a control for the fraction of the population that is non-white, and the number of employees working at the system. The model also includes year and ADD fixed effects.

in unique ways. The regression demonstrates that systems that approve more money for proactive projects are less likely to experience a health-based violation. This effect is unchanged for systems based on the amount of time spent in violation in the previous year. However, the effect of a reactive project is dependent on the amount of time spent in violation. On their own, reactive projects appear to increase the likelihood of a violation. Only those systems that spent more of the prior year in violation and invested in reactive projects are less likely to have a health-based violation. These results further cement that reactive projects and violations are both likely to occur at lower levels of infrastructure quality, but the effect of reactive projects is unclear in this initial analysis. In Section 4, I construct a structural model in which infrastructure quality is the main state variable that drives reactive projects and violations in order to better examine the relationship between investment and quality.

4 Model

I model the system manager's problem as an annual decision to undertake an infrastructure project to improve system quality or to delay investment until the following period. System managers do not know the exact investment costs prior to making an investment decision. However, managers are aware of the distribution of project sizes and use this information to form expectations over the realizations of project costs.

The consequences of noninvestment are driven by the system's infrastructure quality. Each manager knows the current state of their quality and that they face two types of quality uncertainty. The first is an annual shock capturing temporary quality changes that only affect system outcomes in the current period (i.e., storms and chemical spills). This shock and the level of quality at the beginning of the period determine the expected amount of time the system spends in violation and the type of project that is likely to occur if an investment decision is made. The second type of uncertainty is a shock realized at the end of the period that captures lasting unexpected infrastructure changes (i.e., damage due to flooding or unexpected pipe breaks) and leads to uncertainty in the system manager's expectations for future quality levels.

Over time, quality progressively declines. To offset this decline, system managers can increase quality through project investment. If a manager chooses not to invest, the system is more likely to experience a violation and projects are more likely to be reactive. To capture the pressure placed on system managers from providing violating water to consumers, I assume that system managers seek to minimize the weighted costs to consumers from contaminated water exposure against the anticipated costs of an infrastructure project. Below I provide further details on the mechanisms of the model.

4.1 Project Size

In the model, project costs are drawn from pre-estimated distributions. Holding the distributions of project investment size fixed serves two purposes. First, this normalization acts as a budget constraint on manager investment decisions.[14] Data on approved project costs capture the existing limitations and opportunities available to, and perceived by, system managers at the time of approval. By using this data to estimate project size distributions, I incorporate these constraints into the model. For example, the probability of receiving a loan or principal forgiveness for a project is already accounted for by system managers in the observed project sizes. Additionally, any budgetary constraints imposed by political pressure from increasing water costs to consumers and the overhead of planning and developing an infrastructure project are also integrated into existing project costs. Second, the true costs of investment are often difficult to predict and changes to initial estimates can be required to bring projects to completion.[15] By drawing project costs from these distributions, I am able to realistically incorporate the project cost uncertainty faced by system managers.

I construct project size distributions from data in the "anticipated budget" section contained within the WRIS project reports. Using observations for proactive and reactive project costs, I employ maximum likelihood methods to estimate the parameters for two lognormal project size distributions. My analysis confirms that proactive and reactive projects are sourced from statistically different distributions, and that reactive projects are statistically larger than proactive projects.[16] The results of the maximum likelihood estimations reveal a third purpose for holding project distributions fixed. Because reactive projects are more expensive in expectation than proactive projects, system managers have an incentive to invest at higher levels of quality—when a project is more likely to be proactive as opposed to reactive.

4.2 System Quality

The state variable that drives the dynamic discrete choice model is the quality of each water system's physical infrastructure. Infrastructure quality, q_{wt}, is a single measure that captures the functionality of a multitude of drinking water system components and is not observed by the econometrician. I model the distribution of unobserved infrastructure qualities across all systems in 2007, the beginning of my observation period, with a normal distribution, $q_{w0} \sim \mathcal{N}(\mu_q, \sigma_q)$.[17] An $\mathrm{AR}(1)$ process captures the progressive degradation of system quality. Management can combat this quality decay by spending money on infrastructure projects. I estimate different effects for proactive and reactive projects in the model to allow the two types of projects to alter quality in their own way.

I use $i_t = \{0, 1\}$ to indicate whether any project is undertaken in the current period, $y_{rt} = \{0, 1\}$

[14] See Posenau (2022) and Agrawal and Kim (2022) for a detailed analysis of the funding process for local drinking water infrastructure investment. Jerch (2022) demonstrates that localities, when faced with a federal mandate to improve wastewater filtration technology, do not reallocate funds from other areas of their budget but rather increase rates to consumers.

[15] Bajari et al. (2014) study the renegotiation of procurement contracts in a similar infrastructure setting, supporting that exact infrastructure costs are often unknown prior to completion of a project.

[16] Figure E.3 in the appendix displays the fits for these distributions. The results of a one sided Kolmogorov-Smirnov test indicate that the reactive project size distribution first-order stochastically dominates the proactive project size distribution ($D(863, 1454) = 0.086, p < .001$).

[17] In alternate specifications I assume the initial distribution of quality follows a log normal distribution. In this case I recover parameters consistent with the same distribution I estimate when I assume the initial distribution is normal.

to indicate a reactive project, and (k_{pwt}, k_{rwt}) to represent the amount of money spent on proactive and reactive projects respectively.[18]

The law of motion for quality follows the structure below:

$$\begin{aligned} q_{wt+1} &= \alpha^q q_{wt} + i_t\big((1-y_{rt})\alpha^p k_{pwt} + y_{rt}\alpha^r k_{rwt}\big) + \tilde{\epsilon} \\ \alpha^q &\in (0,1) \\ \tilde{\epsilon} &\sim \mathcal{N}(0,1) \end{aligned} \tag{2}$$

where α^q captures the persistence of system quality over time, and $\tilde{\epsilon}$ is realized at the end of the period and captures lasting changes to quality such as damage due to flooding, or an unexpected external grant award. A value of α^q close to one implies that quality degrades slowly and that any changes to quality, here made by project investments, are highly persistent. The multipliers on project investment, (α^p, α^r), reflect the effect of a $1 million expenditure on the level of quality for a proactive and reactive project respectively.

Each period the system receives a temporary shock to system quality, $\epsilon_q \sim$ Type I extreme value, to capture year-specific challenges that the system manager might face. This quality shock introduces uncertainty in both the exact amount of time the system spends in violation and whether a project is proactive or reactive if the system manager decides to invest. I employ an ordered logit model to represent the number of quarters the system spends out of compliance, $y_{vt} = m = \{0,1,2,3,4\}$.[19] As the quality of the system decreases, the expected number of quarters spent in violation increases.

The probability of spending any given amount of time in violation is represented below:

$$Pr(y_{vt} = m) = \begin{cases} 1 - \Lambda\big(q^*_{vm} - q_{wt}\big) & \text{if } m = 0 \\ \Lambda\big(q^*_{vm-1} - q_{wt}\big) - \Lambda\big(q^*_{vm} - q_{wt}\big) & \text{if } m \in \{1,2,3\} \\ \Lambda\big(q^*_{vm-1} - q_{wt}\big) & \text{if } m = 4 \end{cases} \tag{3}$$

where $\Lambda(\cdot)$ denotes the logit function.[20] The thresholds, $\boldsymbol{q_v^*} = (q^*_{v0}, q^*_{v1}, q^*_{v2}, q^*_{v3})$, are decreasing in magnitude and represent the relative values of infrastructure quality at which the ability of the system to provide safe drinking water is diminished and the expected number of quarters in violation increases. For any systems with quality above q^*_{v0}, the expected number of quarters in violation is entirely dictated by the annual quality shock. For systems with quality below this

[18] In the model I restrict system managers to investing in a single infrastructure project in a year. The estimation implications of this assumption are discussed in Section 5.

[19] SDWIS data indicates that violation durations tend to be clustered around month long intervals. I further simplified these observations into quarters of a year spent in violation.

[20] More explicitly, the logit function has the following form:

$$\Lambda(z) = \frac{\exp(z)}{1+\exp(z)}$$

threshold, the probability of spending time in violation increases as quality falls.[21]

The probability that a system undertakes a reactive, as opposed to proactive, project is also determined by current system quality. To mirror the results in Section 3, the probability of a reactive project is higher for lower values of system quality. As quality falls below the reactive threshold, q_r^*, the projects undertaken by system managers are increasingly likely to be reactive. Given the yearly shock to quality, the probability that a system draws a reactive project takes the following form:

$$Pr(y_{rt} = 1) = p_r(q_{wt}, q_r^*) = \Lambda\big(q_r^* - q_{wt}\big) \tag{4}$$

The quality threshold, q_r^*, represents the point at which a reactive project is more likely than a proactive project. This reflects an increase in observable signs that the system will be unable to provide safe drinking water, e.g., breaks and damage, that point to larger problems. As system quality drops, it is increasingly likely that projects are reactive in nature.

Equation (2) indicates that without project investment, system quality declines over time. As system quality declines to lower levels, the probability of spending time in violation and the likelihood of a reactive project increase. The dynamic program estimates the infrastructure quality thresholds $(q_{v0}^*, q_{v1}^*, q_{v2}^*, q_{v3}^*, q_r^*)$ where these outcomes are increasingly likely.[22]

4.3 Consumer Costs from Unsafe Drinking Water

System managers incur costs when they are unable to provide safe drinking water to consumers. I assume that these costs are proportional to estimates for consumer costs from receiving tap water that is below health-based standards. In order to determine the value of consumer costs, I employ a logit-style demand model to estimate consumers' willingness to pay for safe drinking water.

Each week consumers choose between buying bottled water and drinking tap water. I model the indirect utility to consumer i of purchasing bottled water in week t as:

$$u_{i1t} = \gamma^c + \gamma^v v_{it} + \gamma^p p_{1t} + \xi_{1t} + \epsilon_{i1t} \tag{5}$$

where v_{it} is an indicator for consumer exposure to a health-based water quality violation, and p_{1t} is the average price for a gallon of bottled water in week t.[23] Consumer costs from exposure to a water quality violation is captured by γ^v, γ^p captures price disutility, and ξ_{1t} is unobserved by the econometrician and captures a demand shock for bottled water in week t. I assume each person consumes 32 oz of water per day, which equates to 1.75 gallons per person per week.[24] Using this

[21] When estimating the model, I assume the same violation thresholds hold for the duration of the observation period. For further information about changes to the EPA standards and the applicability to the data see Section 2.3.2.

[22] Figure E.4 in the appendix plots simulated data to provide further intuition for the progression of quality and the consequences from noninvestment.

[23] I assume a gallon of tap water is free. As compared to the amount of water used by a consumer for all house-based activities, the relative cost of 1.75 gallons of water per person per week is low. For example, the Muldraugh Water Department currently charges $20.50 for 4,000 gallons of water, which is less than $0.01 per gallon.

[24] There is no definitive rule for how much water consumers need per day. Harvard Health Publishing recommends that healthy people should be drinking four-to-six cups daily (32 - 48 oz). To allow for variation in the types of liquids con-

Table 5: Demand Parameter Estimates and Derived Statistics

Parameter		Estimate	Std. Error
Demand			
Constant	γ^c	-2.256	0.014
Price Parameter	γ^p	-1.007	0.008
Violation Parameter	γ^v	-1.313	0.050
Derived Statistics			
Willingness to Pay	γ^v/γ^p	1.304	
Demand Elasticity	$\gamma^p p_1(1-s_1)$	-1.651	

as a basis for consumption, I construct shares of bottled water and tap water in each county-week observation, s_{1ct} and s_{0ct} respectively.

The results of my estimation are presented in Table 5. Consumers are sensitive to both the price of a gallon of water and exposure to a health-based violation, which is reflected in negative estimates for (γ^p, γ^v) respectively. According to my estimates, each consumer is willing to pay \$2.28 per week (1.75 gallons $\times$ \$1.30) for bottled water to avoid a health-based water quality violation.

4.4 System Manager's Decision Problem

Last period's quality determines the amount of time the system spends in violation during the current period. Therefore, the system manager decides whether or not to invest in a project in the current period in order to affect next period's quality level. System managers know the cost distributions for the possible reactive and proactive projects, but not the true realization of these costs until after deciding to undertake a project.[25]

I model the cost to system managers from providing drinking water that is below health-based standards as proportional to consumer costs, which I construct based on the willingness to pay estimates. To obtain an annual value for these costs, I assume every consumer shops once a week during the year (13 times per quarter) and makes a choice to buy the water needed for that week (1.75 gallons per person) or to consume water from the tap. The annual expected costs to consumers served by community water system w can be represented as:

$$D(q_{wt}, \boldsymbol{q_v^*}) = \big((\hat{\gamma^v}/\hat{\gamma^p}) * 13 * 1.75\big) n_w E\big[y_v(q_w, \boldsymbol{q_v^*})\big] \tag{6}$$

where n_w indicates the average population served by community water system w. During estima-

sumed, I use the lower bound in my estimation. Any deviation from this value is captured in the estimate for the weight system managers place on consumer costs, λ. Sourced from: `https://www.health.harvard.edu/staying-healthy/how-much-water-should-you-drink`

[25]The uncertainty over exact project costs also prevents managers from strategically delaying infrastructure investment in order to receive a "better" project cost draw. Under the timing of the model, delaying the investment decision increases the likelihood of a reactive project which in expectation also increases the anticipated cost of a project. This realistically approximates the repercussions of delaying investment.

tion, I treat all systems identically by holding this value fixed at the average population observed in the data. Every system with the same quality therefore faces the same expected costs of an infrastructure project and of providing unsafe drinking water to consumers. These costs can be conceptualized as normalized per capita values. Based on the estimated parameters, each quarter that the system spends in violation results in consumer costs of approximately $347,000. Equation (6) indicates that consumer costs are increasing in the expected amount of time spent in violation, $E[y_v]$. As quality falls, the amount of expected time in violation increases. This relationship drives system managers to make investments in infrastructure to counter quality decay.

There is an extensive literature examining the true value consumers place on a desired environmental improvement. The primary debate concerns contingent valuation (stated preferences) contrasted with often dissimilar averting behavior (switching to bottled water) estimates of consumer willingness to pay.[26] Along the lines of this argument, I anticipate that system managers are better able to internalize the costs to consumers from exposure to contaminated water than the willingness to pay consumers display when they switch to bottled water. As a result, I estimate a scaling parameter on consumer costs, which I denote λ, in the system manager's optimization problem to capture the anticipated mismatch between consumer valuations and the actual costs considered by system managers. Values of λ less than one indicate that system managers do not fully internalize consumer costs from unsafe drinking water, and values of λ greater than one indicate that system managers are incorporating additional factors into their valuation such that their costs are larger than estimates for consumer willingness to pay.[27]

Let $\boldsymbol{\theta} = (q^*_{v0}, q^*_{v1}, q^*_{v2}, q^*_{v3}, q^*_r, \lambda, \alpha^q, \alpha^p, \alpha^r, \mu_q, \sigma_q)$ denote the set of parameters to be estimated with the dynamic discrete choice model. The distributions for project costs, $k_p \sim \log\mathcal{N}(\mu_p, \sigma_p^2)$, and $k_r \sim \log\mathcal{N}(\mu_r, \sigma_r^2)$ have been pre-estimated using collected project data. I hold fixed the rate at which system managers discount the future at $\beta = 0.95$.

The current period utility takes the following form:

$$u(q_{wt}, i_t; \boldsymbol{\theta}) = \begin{cases} -\lambda D(q_{wt}, \boldsymbol{q_v^*}) & \text{if } i_t = 0 \\ -\lambda D(q_{wt}, \boldsymbol{q_v^*}) - \big(1 - p_r(q_{wt}, q_r^*)\big)E[k_p] - p_r(q_{wt}, q_r^*)E[k_r] & \text{if } i_t = 1 \end{cases} \tag{7}$$

which is dependent on the manager's decision to invest in a project or not. Based on the timing in the model, the cost of pursuing a project in the current period is always greater than choosing to delay investment. System managers make decisions to undertake projects based on the future expected increase in quality from those investments and the change in expected infrastructure costs. As previously demonstrated, delaying investment increases the probability that the system will have to undertake a more expensive reactive project.

Combining the components of the model, the value function for a community water system

[26] The debate is discussed in Bartik (1988), Wu and Huang (2001), and Orgill-Meyer et al. (2018).

[27] Additional forces potentially captured by this weight include the "citizen's voice" phenomenon discussed in Brooks and Liscow (2023), who attribute growing infrastructure costs to increased public influence in government decision making or political pressure from other branches of the government for non-compliance with health-based standards.

manager has the following form. I introduce an additional set of state variables, $\big(\epsilon(0), \epsilon(1)\big)$ for other choice-specific states that are only observable to the system manager. For notational simplicity, I omit time and water system subscripts.

$$V\big(q, \boldsymbol{\epsilon}; \boldsymbol{\theta}\big) = \max_{i \in \{0,1\}} \Big\{ u(q, i; \boldsymbol{\theta}) + \epsilon(i) + \beta EV\big(q', \boldsymbol{\epsilon'} | i, q; \boldsymbol{\theta}\big) \Big\} \tag{8}$$

5 Estimation

5.1 Methodology

To estimate the model parameters, I use observations of system manager decisions to undertake infrastructure projects, realizations of the type of infrastructure project, and the amount of time systems spend in violation to construct a maximum likelihood function. The behavior of system managers is driven by three state variables that are unobserved by the econometrician: infrastructure quality and two choice-specific shocks. The estimation of models with unobserved choice-specific shocks is covered extensively in the literature, and in keeping with this work I make some preliminary simplifying assumptions about the relationships between the error terms and system quality.[28] I assume that the error terms $\epsilon(i)$ are independent in every period and follow a normalized Gumbel distribution with mean zero and variance $\pi^2/6$. With these assumptions I am able to use fixed-point methods to construct conditional choice probabilities for project investment decisions as a function of infrastructure quality. The main challenge for estimation is that quality is unobserved by the econometrician and depends on the value of quality in the prior period. This history dependence makes estimation of the model computationally intensive as all possible values of quality across 13 periods of observation need to be considered in the evaluation of the maximum likelihood function.

Reich (2018) proposes a new method called recursive likelihood integration (RLI) that applies recursive estimation techniques to reduce computational complexity in models with serially correlated continuous unobserved states. RLI incorporates iterative interpolations and approximations in conjunction with backward induction to solve a likelihood function over possible values of an unobserved state variable.

Data on the timing of project investments, i_t, the number of quarters spent in violation, y_{vt}, realizations of reactive projects, y_{rt}, and project costs, (k_{pt}, k_{rt}) identify the parameters of interest. The likelihood function for a single system can be represented as follows.

$$\begin{aligned} L_T(\boldsymbol{\theta}) \quad = \quad & \int ... \int \prod_{t=1}^{T} \Big(p_{i|q}(i_t | q_t; \boldsymbol{\theta}) p_{y_r|q}(y_{rt} | q_t, i_t; \boldsymbol{\theta}) p_{y_v|q}(y_{vt} | q_t; \boldsymbol{\theta}) \\ & \times \quad p_q(q_t | i_{t-1}, y_{rt-1}, k_{pt-1}, k_{rt-1}, q_{t-1}; \boldsymbol{\theta}) \Big) \\ & \times \quad p_{i|q}(i_0 | q_0; \boldsymbol{\theta}) p_{y_r|q}(y_{r0} | q_0, i_0; \boldsymbol{\theta}) p_{y_v|q}(y_{v0} | q_0; \boldsymbol{\theta}) p_q(q_0; \boldsymbol{\theta}) dq_T dq_{T-1} ... dq_0 \end{aligned} \tag{9}$$

[28] A survey of dynamic discrete choice models and methods of estimation is covered in Aguirregabiria and Mira (2010).

The goal of recursive likelihood integration is to reduce the above equation to T one-dimensional integrals, which are easier to compute, and to apply techniques of backward induction over the history of qualities in the period of observation to solve the maximum likelihood function. To achieve this goal, I follow Reich (2018) and define the following recurrence relation.

$$g_t(\tilde{q}) = \begin{cases} 1 & \text{if } t > T \\ \int p_{i|k_p,q}(i_t|\tilde{q}';\boldsymbol{\theta}) p_{y_r|q}(y_{rt}|\tilde{q}', i_t;\boldsymbol{\theta}) p_{y_v|q}(y_{vt}|\tilde{q}';\boldsymbol{\theta}) & \text{if } t \leq T \\ \quad \times\, p_q(\tilde{q}'|i_{t-1}, y_{rt-1}, k_{pt-1}, k_{rt-1}, \tilde{q};\boldsymbol{\theta}) g_{t+1}(\tilde{q}') d\tilde{q}' & \end{cases} \tag{10}$$

I employ the above recurrence relation to recursively determine the most likely parameters that generated the observed data, given all the possible paths quality could have taken from the first period to the last period. For each parameter guess, I use equation (10) to build up the likelihood of the observed data over all the possible values of q in each time period. To construct the likelihood, I first discretize quality by establishing a set of grid points for evaluation. I then initialize the recurrence relation to 1 in period T, setting the probability of being at any possible point across the grid of q values in period $T+1$ to be equivalent because $T+1$ outcomes are not observed. Then I calculate, for each grid point of q, the probability that the system quality is that value at period T, given the possible values system quality could have been at $T-1$ and the realizations of quality consequences in period T. Using backward induction, I repeat this process until $t = 1$ to establish the likelihood of a quality path from period $t = 1$ to $t = T$ at each of the grid points of q. At period $t = 0$, I make a final approximation over the possible initial values of quality based on the distribution of q_0, and the observations of project decisions, violations, and reactive projects in the initial period. Additional details on the RLI process can be found in Appendix B.

5.2 Identification

Identification of the parameters is driven by data on the timing, type, and amount of project investments, data on the number of quarters spent in violation, and distributional and functional form assumptions for the progression of quality. Figure 1 depicts the observed data for four different systems. Quality is assumed to be, on average, decreasing over time with increases in quality occurring when managers invest in projects. Although quality is unobserved, Figure 1 can be thought of as a representation of the consequences of an ungraphed quality level that determines when the system manager invests, the type of investment made, and the quarters spent in violation. Data on these consequences and variation in behaviors and outcomes across systems help to identify the parameters dictating the progression of system quality, the thresholds for violation, and the threshold for reactive projects.

To understand how the parameters are identified, first consider a model where the quality of each system is observed by the econometrician. With values of quality observed and data on the number of quarters spent in violation, the thresholds for violation, $\boldsymbol{q_v^*}$, are quickly pinned down. The threshold for reactive projects, q_r^*, is derived from quality levels at the time of investment

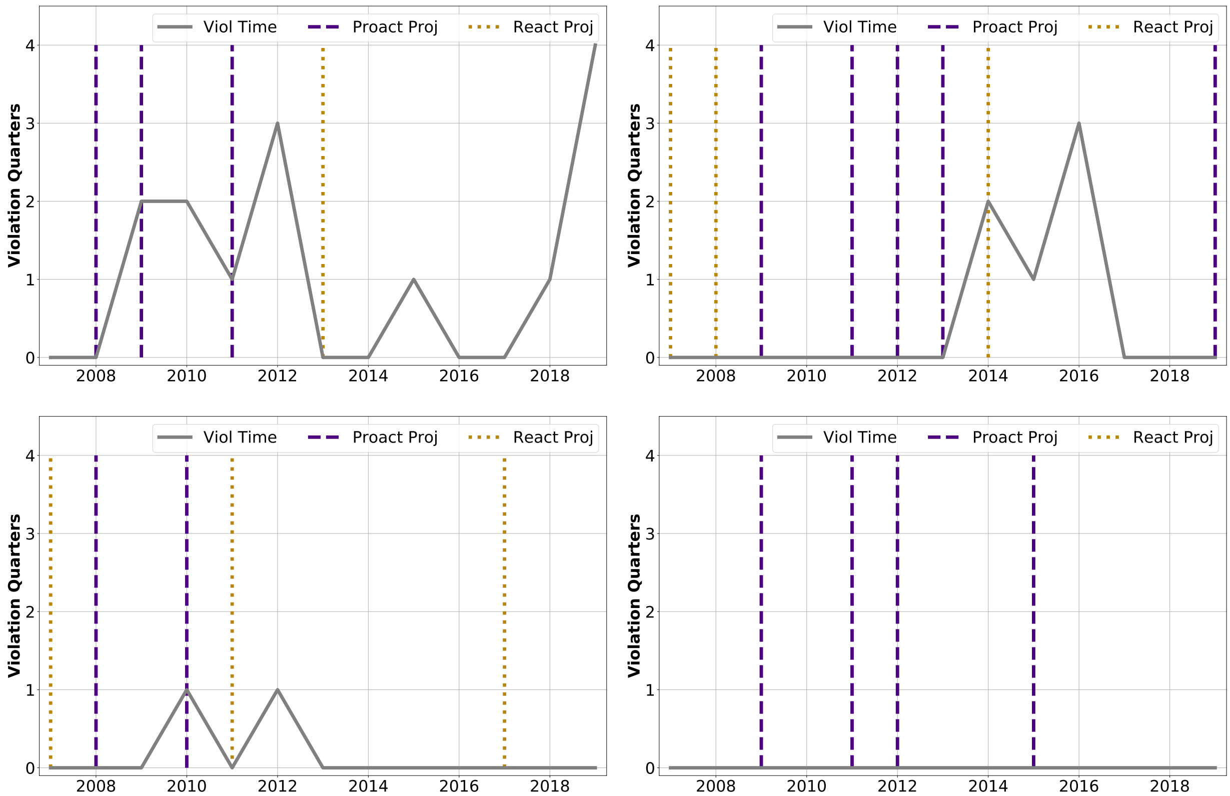

Figure 1: Timing of Projects and Violations

Notes: The panels above portray observed water system activities for four different systems. Gold tightly dotted vertical lines indicate a reactive project investment, purple dashed lines indicate a proactive project. Solid gray lines indicate the number of quarters spent in violation.

and realizations of project types. An OLS regression of current quality on lagged values of quality and observed project costs identify the parameters describing the law of motion for q: the $\mathrm{AR}(1)$ parameter, α^q, and the change in quality from a million dollar investment in each type of project, (α^p, α^r). The weight that system managers place on the cost of unsafe drinking water exposure to consumers, λ is identified by the timing of infrastructure project investment, i.e., when the expected cost of investment is less than the cost incurred by inaction.

When considering a model where quality is unobserved, many of the same identifying conditions hold, but quality no longer serves as a grounding foundation for the identification of these parameters. For instance, the identification of the thresholds described above depends on the observed value of quality when systems enter into a state of violation or undertake a reactive project. Without observations of quality, the parameters in the model are no longer fully identified, only the relative levels of the parameters. To remedy this identification issue, I make a few normalizations.

First, I hold the mean of the initial distribution of quality fixed, i.e., $\mu_q = 50$. With this normalization, the model is able to "anchor" the parameters based on the average initial quality level, and all other parameters are interpreted in this context.[29] Second, I assume that α^q, the rate of quality

[29] Simulations reveal that similar likelihood values can be obtained using alternate values for this parameter, pointing

deterioration, is 0.99, which ties average quality to the penalties from non-investment. I fix the $\mathrm{AR}(1)$ parameter close to 1 in order to mimic a setting where quality deteriorates slowly over time. Similar to the discount rate, the rate of quality deterioration is not uniquely identified in the model. Because I estimate the thresholds for consequences from noninvestment and the increase in quality from a million dollar investment, the observed data is not sufficient to pinpoint a specific value for α^q. For example, α^q close to zero paired with large values for (α^p, α^r), the effect of investment on quality, and low violation and reactive thresholds can justify the same gaps in infrastructure investment decisions as a value of α^q close to one paired with small values for (α^p, α^r), and high violation and reactive thresholds.

With the above normalizations, the rest of the parameters in the model are identified from a combination of these values and observed data. As depicted in Figure 1, the "gaps" between project investment and time spent in violation help to pin down the parameters of interest. Because the rate of quality decline is fixed, the initial violation threshold is determined by the number of periods that pass before a system spends time in violation. The additional violation thresholds are then identified by how many periods pass without investment before the amount of time in violation increases. Similarly, realizations of reactive projects and the number of periods between activity isolate the reactive project threshold, q_r.

The estimated effect of investment on quality is linked to the weight placed on consumer exposure costs. Projects are only undertaken when the expected expense of investing in an infrastructure project is outweighed by the expected benefit from investment. The magnitude of investment benefit is defined by the increase in quality and the weight placed on the costs from entering into a state of violation. Variation in realized expenditure types and amounts separately identify the effect of proactive and reactive projects. The timing of investment then identifies the multiplier on consumer costs from violations.[30] Lastly, the spread of the initial quality distribution, σ_q, follows from the initial distribution of systems in violation and the overlapping distributions of the highest probability q_0 values for each water system holding the other normalizations fixed.

5.3 Estimation Results and Analysis

5.3.1 Parameter Estimates and Interpretation

I estimate the structural parameters of the model using data on the project decisions and violation outcomes of 353 community water systems from 2007-2019.[31] The point estimates from the dynamic program are presented in Table 6. As infrastructure quality declines, system managers are less likely to be able to provide consumers with safe drinking water. Estimates of the violation

to an identification issue. See Appendix C for further details.

[30] Note that I also do not estimate a multiplier on the expected cost of investment. Because quality is unobserved and I estimate the effect of realized investment on quality, the model can only recover either this multiplier or the weight managers place on consumer costs.

[31] In the event that a system approves multiple projects in a year, I pool observations by project type into a single investment. In approximately 3% of observations both reactive and proactive projects are approved in the same year. In these cases, I keep the project type with greater total expenditure.

Table 6: Dynamic Model Parameter Estimates

Parameter		Estimate	Std. Error
Delay Consequences			
Consumer Cost Weight	λ	9.522	1.014
Reactive Project Threshold	q_r^*	47.456	0.389
1Q Violation Threshold	q_{v0}^*	44.776	0.275
2Q Violation Threshold	q_{v1}^*	43.633	0.269
3Q Violation Threshold	q_{v2}^*	42.504	0.242
4Q Violation Threshold	q_{v3}^*	41.635	0.225
Project Effect on Quality			
Proactive Investment	α^p	0.061	0.006
Reactive Investment	α^r	0.038	0.005
Initial Quality Distribution			
Initial Quality Variance	σ_q^2	1.248	0.919
Fixed Parameters			
Initial Quality Average	μ_q	50	
Deterioration Rate	α^q	0.99	

thresholds indicate that a system starting at average initial quality in 2007 can expect to reach violating quality levels within 11 years if the manager undertakes no infrastructure projects. With continued noninvestment, the average system crosses the threshold for spending the entire year in violation within 19 years. Assuming that spending the entire year in violation of health-based standards signals the end of a system's usable life, my analysis is consistent with EPA estimates.[32]

The relative benefit from a project depends on the level of quality at the time of investment, and how close quality is to the violation thresholds. For instance, at a quality of two, investing in a project on average increases quality by almost 6%. However, for a quality level that low, there is no benefit from an investment because the number of quarters the system spends in violation remains unchanged. Figure 2 plots the cumulative benefit of different levels of project investment depending on the value of quality at the time of investment. For both types of projects, the most advantageous time to invest is when quality is close to the violation thresholds, i.e., when the system is likely to enter into violation in the following period. At this point, an investment decreases the amount of time the system spends in violation in the following period by the maximal amount. If system managers invest in infrastructure prematurely, the investment is less effective at reducing the time spent in violation due to the progressive degradation of system quality. For example, an average proactive project investment made at a quality level of 55 reduces the amount of time

[32]The Environmental Protection Agency (2002) report "The Clean Water and Drinking Water Infrastructure Gap Analysis", states that water treatment plants typically have a useful life of 20-50 years. Given that most systems are not new in 2007, the results indicate that holding fixed the initial value of quality and the rate of deterioration do not hinder the model's ability to match related external estimates for systems' useable life.

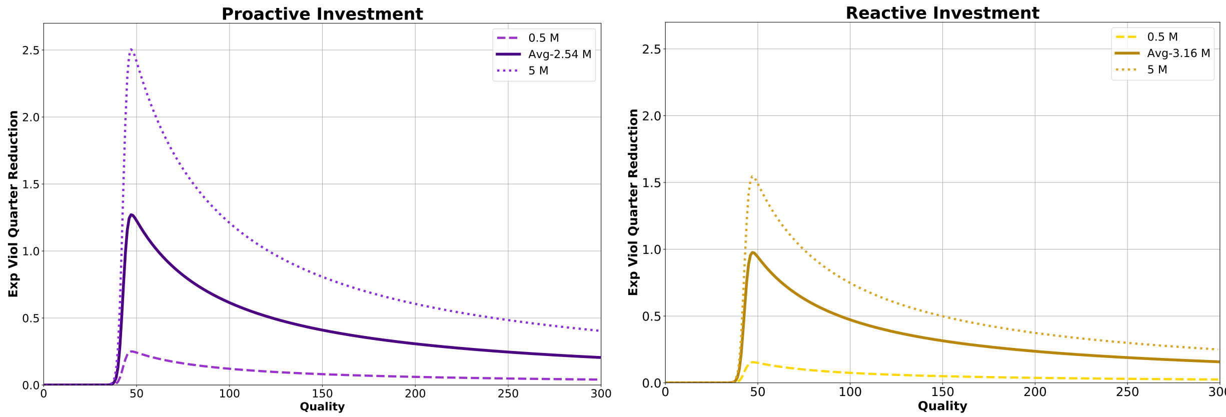

Figure 2: Value of Project Investment

Notes: The panels display the cumulative expected reduction in violation quarters for different levels of investment based on the value of quality when the investment is undertaken.

spent in violation by approximately 88% of the same investment made just before entering into a state of violation, at a quality level of 47.

System managers also face an increased probability of having to undertake a more expensive reactive project as quality falls to lower levels. The parameter estimates for the first threshold for violation and the threshold for a reactive project dictate which of these consequences managers are likely to encounter first. The reactive project threshold is higher than the first threshold for time spent in violation, indicating that managers observe signs of quality decline before experiencing a violation. A system with average quality in 2007 is more likely to have to undertake a reactive, as opposed to a proactive, project within six years. The threshold for a reactive project, 47.5, is also very close to the point at which an investment will have the largest expected reduction in the amount of time the system spends in violation.

Parameter estimates for the increase in quality from million-dollar investments provide further context for the tradeoffs between proactive and reactive investment. The average proactive project costs \$2.54 million, and the average reactive project costs \$3.16 million. As a result, for both projects to result in the same increase in quality, α^p/α^r must be at least as large as 1.24. I find that this ratio is 1.62, indicating that managers prefer proactive investment not only because proactive investment is cheaper but also because it is more efficient—proactive projects improve infrastructure quality more than reactive projects. Therefore, a manager with infinite funds and no uncertainty in infrastructure quality would maintain system quality around 47.5. At this level of quality, managers can invest in efficient proactive investment, employing a just-in-time strategy to prevent violations and maximize the effectiveness of every dollar spent. Under current conditions, managers are unable to maintain this level of quality. Due to the size of infrastructure investments, the average proactive project can only reduce the amount of time spent in violation by 1.27 quarters, and the average reactive project by 0.98 quarters.

The multiplier on the exposure costs to consumers, λ, is determined by the investment behavior observed in the data and the estimated effectiveness of projects. To justify the observed pattern of investments, λ is estimated to be 9.52, which is well above one. At a value of one, system managers weight the estimated costs to their consumers from receiving contaminated water at the cost consumers display through their aversion behavior. Here λ is estimated to be greater than consumers' displayed willingness to pay, in line with expectations that system managers are more averse to violations. There are many possible explanations for this estimate of λ. One possibility is that consumers make poor evaluations of the health risks associated with contaminated water. Another explanation is that systems imperfectly inform consumers of violations, and therefore the estimated aversion responses are actually lower bounds on consumer willingness to pay for safe drinking water.[33] I construct willingness to pay estimates from Nielsen data, and primarily use WRIS data to estimate the parameters of the dynamic model. The estimate for the weight placed on consumer costs is consistent with predictions for system manager aversion to providing contaminated water, reflecting that the Nielsen data lends validity to my estimation results.

To assess how well the model fits the data, I simulate 100 scenarios to approximate the behavior of system managers. In each simulation, I draw initial quality from the estimated distribution and then allow systems to make investment decisions. Overall, the model is able to match the trends in the data. Figure 3 compares the average results across all simulations to the observed behavior of systems in the data. In the top four panels, gray lines indicate the annual number of systems that spend 1-4 quarters in violation of a health-based standard. The blue dashed lines indicate the average predicted number of systems in violation as a result of the simulations. Most systems spend no time in violation, indicating that during my period of observation, systems typically have a water quality level above the standards set by the EPA. Over time, more systems spend time in violation as infrastructure quality decreases.

The bottom two panels in Figure 3 compare the true and predicted investment behaviors of system managers. The left panel plots the number of systems with proactive projects in every year and the right panel plots the number of reactive projects. The observed data are summarized by the gray lines and the dashed colored lines portray the predictions. Over time, the number of proactive projects declines. As the average quality of all systems falls to lower levels, the number of systems investing in projects is decreasing, and a larger proportion of projects are reactive. Reactive projects are estimated to be less effective than proactive projects, so as quality declines, the average project becomes both more expensive and less effective. Therefore, the incentive to undertake any project is slowly declining, even as the expected amount of time spent in violation is increasing.

5.3.2 Long Term Model Implications

The parameter estimates also provide sufficient information to simulate the long term state of water systems. Figure 4 plots the histogram of system qualities, the future predicted project trends, and expected quarters each system will spend in violation. The top panel plots simulated quality

[33] Collier et al. (2021) estimate direct healthcare costs for waterborne diseases at $3.33 billion annually. Marcus (2022) finds that individuals can successfully avoid exposure to drinking water contaminants only when notification occurs very close to exposure.

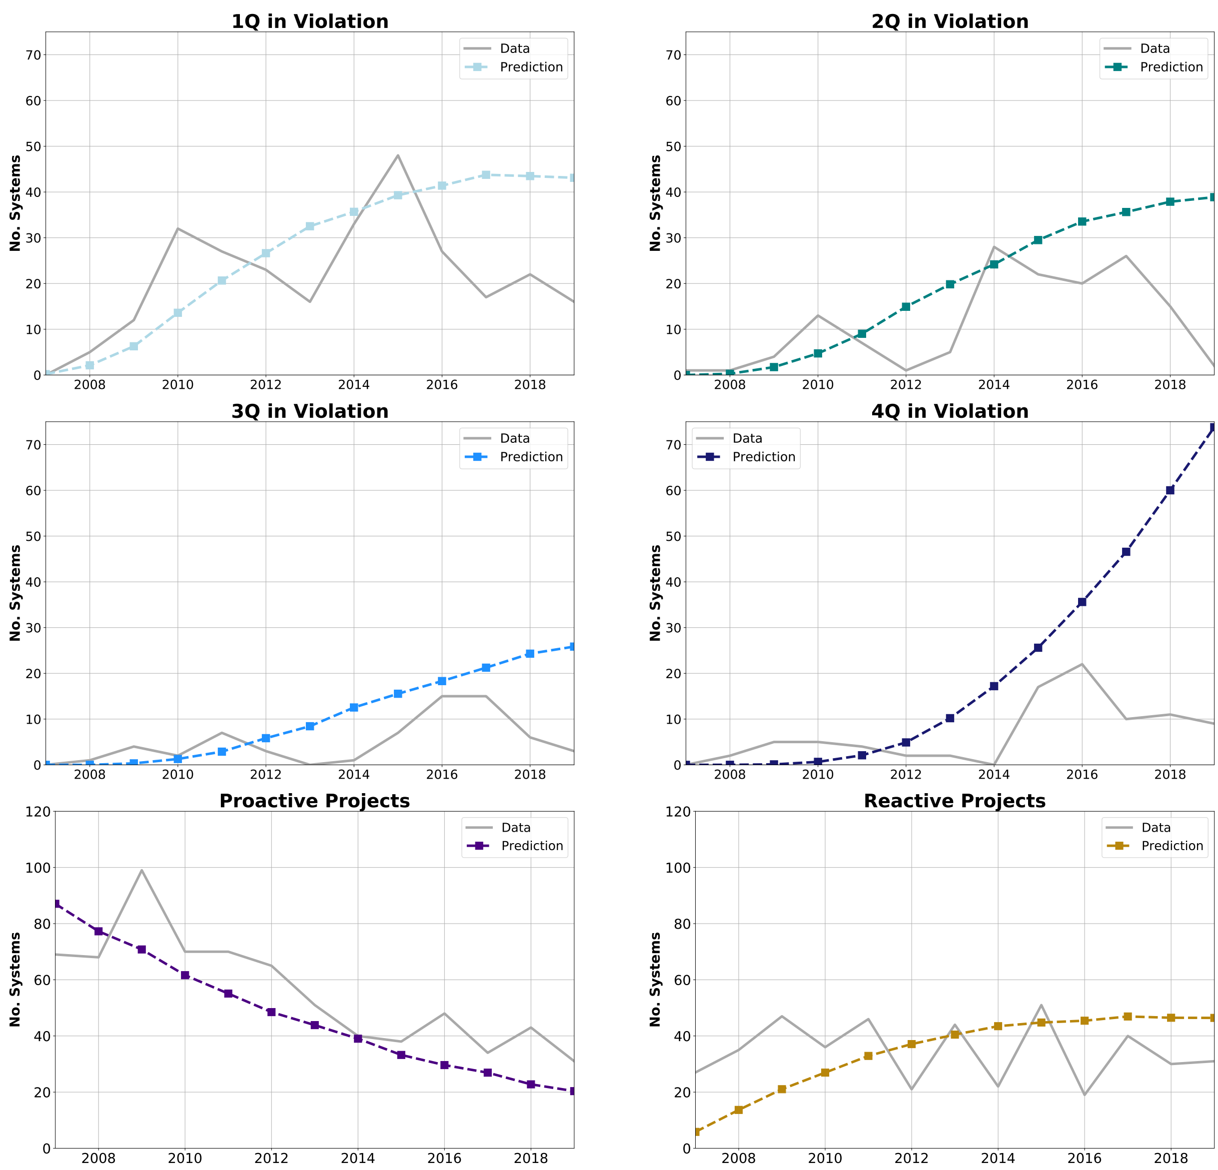

Figure 3: Model Fits

Notes: In each panel, the gray lines indicate the number of systems observed in the data and the dashed lines show the average number of predicted systems calculated from 100 simulations generated from parameter estimates.

distributions in 2010, 2030, and 2050. Over time, the quality distribution becomes increasingly dispersed and steadily shifts downward to increasingly lower levels. To ease interpretation, recall that the average proactive and reactive investments, at their maximal effectiveness, decrease the time spent in violation by roughly one quarter (Figure 2). At the current rate and magnitude of investment, projects are unable to prevent systems from slowly declining to lower and lower levels of quality. The bottom left panel portrays the resulting number and types of projects undertaken by system managers as quality declines. As the average quality level across all systems falls, systems shift to only undertaking reactive projects, and the number of projects progressively declines.

The bottom right panel portrays the predicted average quarters systems will spend in violation over the next 50 years at the current model estimates. The increasing trend in average expected quarters in violation is driven by the progressive change in the quality distribution. As more of the quality distribution falls below the violation thresholds (indicated by the dashed vertical lines in the top histogram), more systems spend longer fractions of the year in violation of a health-based standard. Without policy intervention, more systems will violate these standards and spend increasing amounts of time in violation.

6 Counterfactual Simulations and Policy Implications

The results presented in Section 5 indicate that the rate of project adoption and the size of infrastructure projects are currently too low to prevent system decline into an extended state of violation. In this section, I examine outcomes under alternative policies and determine the ability of these policies to combat quality decline. I first investigate the efficacy of policies designed to increase the number of projects. I find that on their own, these policies are insufficient to produce substantial changes in violation predictions. I then explore a policy that increases the size of projects and discover that this leads to systems spending more time in violation. In my final analysis, I examine policies that include incentives to increase the size and rate of investment. I find that a policy that combines moderate proactive incentives with significant reactive incentives can successfully motivate system managers to reduce future violations to very low levels.

6.1 Initial Policy Simulations

In this section, I first consider the ability of penalties and subsidies to induce different long term violation outcomes holding fixed the resource constraints faced by system managers. I start by simulating scenarios where the state or federal government penalizes violating water systems proportional to consumer exposure costs. The EPA currently uses penalties to incentivize compliance behavior, but penalties are generally applied only to the most egregious offenders.[34]

Under existing conditions, the model predicts that the average system will spend 126 quarters, just over three decades, in violation of health-based standards from 2007-2057. When instituting a

[34] See Blundell et al. (2020) for an analysis of the dynamic application of penalties by the EPA.

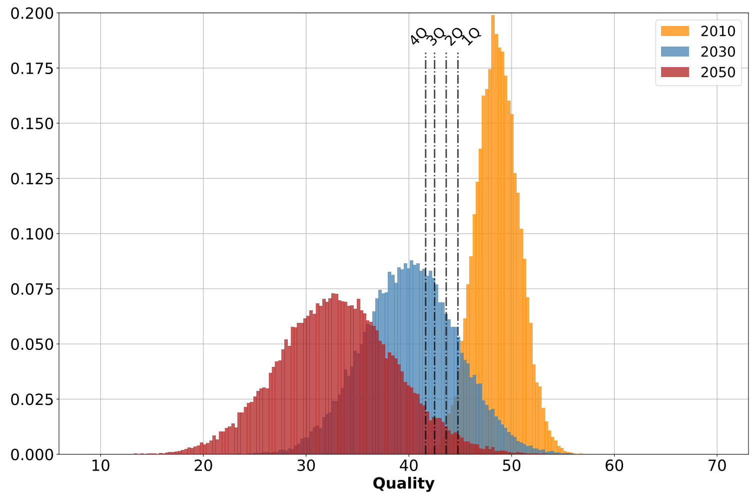

(a) Simulated Quality Distribution

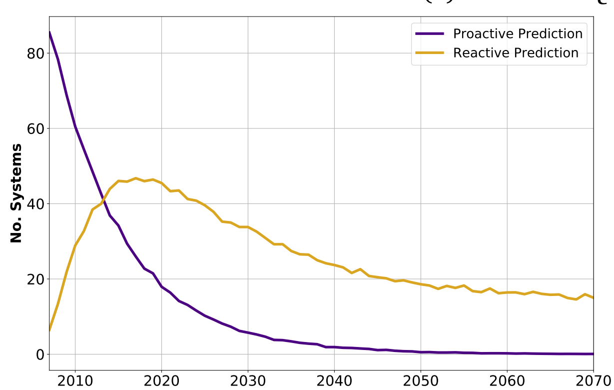

(b) Long Term Project Trends

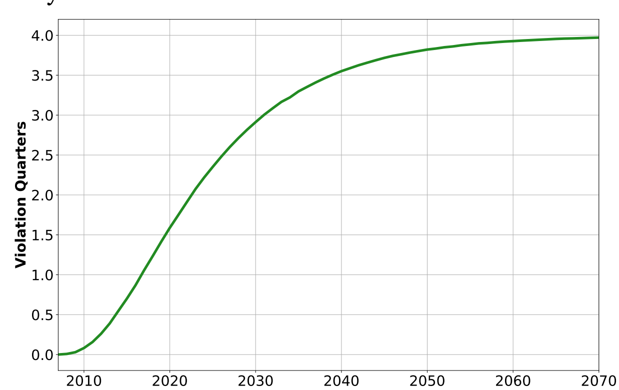

(c) Expected Quarters in Violation

Figure 4: Long Term Implications

Notes: The top histogram portrays the distribution of quality for 100 simulations of 353 systems in 2010, 2030, and 2050. Vertical lines indicate the estimated violation thresholds. The bottom left panel displays the average predicted number of projects by type from 2007-2070. The bottom right panel displays the predicted average quarters each system spends in violation from 2007-2070.

small penalty for violation, five times consumer exposure costs, the model predicts that the amount of time spent in violation over this period will decrease by only one year. Increasing the penalty to 50 times consumer exposure costs reduces violation time by 4.7 years. There is little increased benefit from multiplicative penalties above 50. At these penalty levels, any increase in the cost of spending time in violation no longer serves as an effective incentive to induce system managers to make additional investments. Due to the relatively small average change in quality from an infrastructure project, the manager is not incentivized to make further investments even under severe penalties. This simulation indicates that, holding all else fixed, there is an upper bound on the ability of penalties imposed by the EPA to increase the number of infrastructure projects.

Another way to achieve a similar goal is to institute a policy that decreases the cost to water utilities from investing in infrastructure. In the second set of simulations, I simulate scenarios where projects are fully subsidized and keep the size of investments the same. I find that even with fully subsidized projects, the amount of time spent in violation from 2007-2057 is reduced by only 3.7 years. Similar to the effect of increasing penalties for violations, eliminating project costs to incentivize investment is also unable to prevent system quality decline.

Lastly, I examine policies that increase the size of infrastructure projects, relaxing the budget constraint for system managers but holding fixed system managers' disutility from investment. I implement this policy as an additive increase to the existing project size distributions. In 1996, Congress amended the Safe Drinking Water Act to allow states to specifically set aside funds to support systems with their non-infrastructure needs. Simulations that increase project size can be conceptualized as expanded federal support through this program for the planning, development, and execution of projects as opposed to direct cost subsidies.

Under these conditions, a single project increases infrastructure quality by a larger amount which, in turn, reduces the expected time spent in violation. However, system managers make *fewer* project investment decisions under these conditions. The decrease in project adoption is due to the parameters governing the relative costs and benefits from investment. Although larger projects lead to less time spent in violation, this decrease is not sufficiently offset by the increased investment cost. With this policy, systems spend an additional year, on average, in violation as compared to baseline conditions, indicating that this policy would be counterproductive to reducing the amount of time that systems spend in violation. The results in this section suggest that each of these policies on their own are insufficient to significantly alter the long term predictions for utilities.[35] In the next section, I examine policies that combine increased project size with government subsidies.

6.2 Targeted Policy Simulations

This section explores the benefits of policies that target investment types by first simulating violation outcomes under policies that only promote proactive projects, then policies that only promote reactive projects. I "promote" a project type by completely subsidizing the costs associated with that type of project and by increasing project size. The Drinking Water State Revolving Fund (DWSRF) is an existing government assistance program designed to financially aid water systems in meeting federal health-based water quality standards. As part of this program, states are able to set their own criteria for allocating funds to local systems. Policies that promote a type of investment could be implemented through this or similar programs. However, the purpose of the following simulations is to demonstrate possible outcomes for alternate policies, abstracting from the complexity that would be involved in implementation.

Table 7 presents the violation and project results over 2007-2057 for the baseline estimated conditions and each of the counterfactual policies that I analyze.[36] Comparing outcomes, the proactive-promoting policy is more effective at reducing the average amount of time spent in viola-

[35]Figure E.5 in the appendix depicts how each of these policies alter the average time spent in violation from the baseline prediction.

[36]Figure E.6 in the appendix graphs the average expected quarters in violation under different project size increases for proactive and reactive investments. I focus on the policies that increase project size by $30 million because these policies both achieve very low average violation levels.

Table 7: Counterfactual Policy Comparison, Over 50 Years

	Baseline	Fixed Proactive	Fixed Reactive	Estimated Incentives
Project Size Increases				
Proactive ($ M)	—	30	—	11.84
Reactive ($ M)	—	—	30	110.6
Policy Outcomes				
Violations				
Average Years in Violation	31.43	0.96	1.16	0.03
Average Years in Violation, Violating Systems	31.51	24.01	1.32	0.31
Systems with Violations	99.8%	4.0%	88.1%	8.2%
Projects (Per System)				
Average Proactive Projects	2.184	13.78	7.687	25.55
Average Reactive Projects	4.160	0.654	16.53	1.745

tion and reduces the number of systems that violate health-based standards by the largest amount. Under this policy, the average years spent in violation drops from 31 to less than one and only four percent of systems violate as compared to 99%. However, violating systems spend an average of 24 years, or nearly half the simulated amount of time, out of compliance. The conditions generated by this policy are predominantly successful at reducing future violation time but concentrate violation risk into a few systems that are left vulnerable to catastrophic failure.

By comparison, the reactive-promoting policy reduces the unconditional amount of time spent in violation by a smaller amount, but the difference between the conditional and unconditional time spent in violation is very close. This policy leads to nearly every system (88%) violating health-based standards, but the violation outcomes are less extreme. The average violating system can expect to spend only 1.3 years in violation over the next 50 years. These simulations demonstrate the advantages of the two types of investment incentives: promoting proactive investment is effective but reactive promotion more evenly distributes risk.

In the last column of Table 7, I present the results under a policy that combines proactive and reactive project size increases with project cost subsidies. I calculate these increases by redistributing the funds needed to obtain the reactive-promoting policy to a policy that achieves the minimum average amount of time violating systems spend out of compliance over the simulated 50-year period. The purpose of this exercise is to determine if there is a set of combined incentives that can leverage the advantages of each investment type to improve violation outcomes. Note that this is not an optimal policy, but rather an analysis of how a fixed amount of funding can be distributed to attain the most effective and equitable violation outcomes.

I find that a policy that increases proactive projects by $11.8 million and reactive projects by

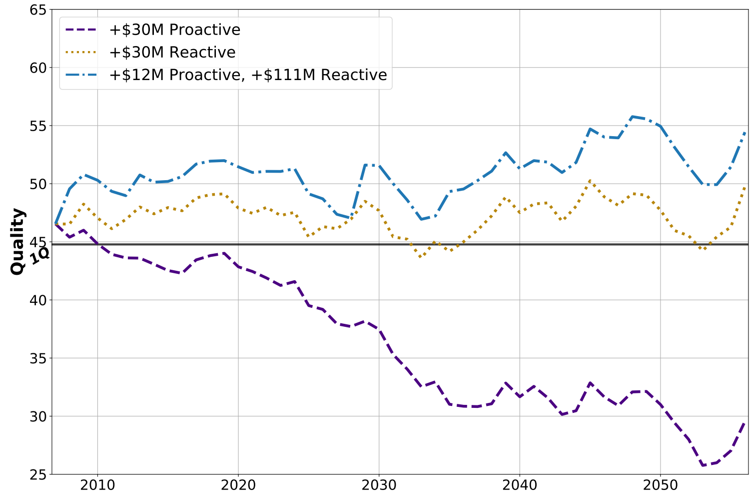

Figure 5: System Quality Outcomes, Counterfactual Policies
Notes: The figure plots simulated quality levels for a sample system over 2007-2057 under the three policies analyzed.

$110.6 million achieves this goal. With these incentives, systems primarily invest proactively (on average 26 times over the 50-year period) and only occasionally undertake reactive projects (on average 1.7 times during the period). Parameter estimates foreshadow this result. As discussed in Section 5.3, proactive projects are cheaper and more efficient. As a result, managers require minimal incentives to invest proactively. The increase in proactive project size is primarily to increase the average quality improvement from investment. On the other hand, an increase of $110 million in reactive project size is required to provide managers with sufficient funds to recover from undesirably low levels of infrastructure quality. Under simulated conditions, the average reactive project improves quality by 4.28, which when implemented just as the system crosses the threshold for a reactive project, at 47.5, is sufficient to prevent the system from violating for the next 15 years without further investment. As a result, disadvantaged systems that start with low initial quality levels or systems that receive particularly negative quality shocks are able to recover and return to higher infrastructure quality levels through reactive investment.

Figure 5 plots the simulated quality paths for an example system under each of the three counterfactual policies presented in Table 7. To demonstrate the full impact of the different policies, I selected a system that experiences violations under the proactive-promoting policy. The example system starts with a low level of infrastructure quality in the initial period. Under the proactive-promoting policy, the manager is unable to recover from this disadvantaged beginning, and the system spends most of the simulated years in violation of health-based standards. In the simulation where reactive projects are promoted, the manager is able to bring initial infrastructure quality up to acceptable levels. However, because the reactive threshold is very close to the threshold for violation the manager maintains a precarious average quality and occasionally falls out of compliance. Finally, the policy that redistributes reactive funds to both proactive and reactive incentives enables the manager to repair the system's initial infrastructure quality and then maintain quality levels such that the system does not experience violations for the duration of the simulation.

7 Conclusion

In this paper, I provide a framework for analyzing infrastructure investment decisions and the effectiveness of proactive and reactive expenditures. I use empirical findings from a new dataset to motivate and estimate a dynamic discrete choice model of water system manager infrastructure investment. The results indicate that smaller proactive investments are required to have the same increase in quality as larger reactive expenditures, implying that proactive investments are more efficient. Simulations demonstrate that at the current levels of investment, water systems will increasingly violate health-based standards in the future. I find that policies targeting system manager investment have different risks. Increased and subsidized proactive investment can enable most managers to invest efficiently but leaves some systems vulnerable to unrecoverable emergencies. Increased and subsidized reactive investment leads to a higher level of violation risk incurred by all systems but reduces disparities between systems. A policy that promotes moderate proactive projects and large reactive projects results in predominantly efficient expenditures and provides a safety net against low states of infrastructure quality.

Although my results align with patterns observed in the data and factors that are attributed to recent infrastructure failures in Jackson, Mississippi and Flint, Michigan, my approach is limited in certain respects. First, I do not have measures of infrastructure quality. I overcome this challenge by implementing new methods to capture the influence of a persistent, unobserved state. I am therefore able to estimate the model and perform counterfactual analysis but because quality remains unobserved, the model does not support implementable recommendations for specific systems. Second, I do not model the complicated relationship between locally managed systems and the state which reviews and funds projects. The state could play an important role in balancing the resources available for infrastructure improvement with the needs of individual systems. Future research into this relationship and the motivations of the state could provide a more complete picture of the entire investment cycle.

My research has contemporary policy implications. The EPA's 2015 Drinking Water Infrastructure Needs Survey indicates that system managers anticipate nearly $500 billion will be required over the next 20 years to maintain properly functioning drinking water infrastructure. Partially to address these and other anticipated infrastructure needs, the Infrastructure Investment and Jobs Act was signed into law at the end of 2021. The Act provides $550 billion in federal support for a host of public infrastructure systems including: passenger rail, highway, drinking and wastewater, high-speed internet, power generation, and electric vehicle charging stations. Much of this federal assistance will be used to upgrade existing systems, yet the difference between reactive and proactive spending is neither mentioned nor considered in the Act. My results indicate that proactive investment is more efficient than reactive investment but that policies promoting proactive investment alone lead to unequal outcomes across all systems. A policy providing incentives for proactive investment paired with reactive investment incentives to correct for unexpected disasters enables all managers to maintain functional infrastructure.

References

ABITO, J. M. (2019): “Measuring the Welfare Gains from Optimal Incentive Regulation,” *The Review of Economic Studies*, 87, 2019–2048.

AGRAWAL, A. K. AND D. KIM (2022): “Municipal Bond Insurance and Public Infrastructure: Evidence from Drinking Water,” *Draft*.

AGUIRREGABIRIA, V. AND P. MIRA (2010): “Dynamic discrete choice structural models: A survey,” *Journal of Econometrics*, 156, 38–67.

ALLAIRE, M., T. MACKAY, S. ZHENG, AND U. LALL (2019): “Detecting community response to water quality violations using bottled water sales,” *Proceedings of the National Academy of Sciences of the United States of America*, 116, 20917–20922.

ALLAIRE, M., H. WU, AND U. LALL (2018): “National trends in drinking water quality violations,” *Proceedings of the National Academy of Sciences of the United States of America*, 115, 2078–2083.

ARCIDIACONO, P. AND R. A. MILLER (2011): “Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity,” *Econometrica*, 79, 1823–1867.

BAJARI, P., S. HOUGHTON, AND S. TADELIS (2014): “Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs,” *American Economic Review*, 104, 1288–1319.

BARTIK, T. J. (1988): “Evaluating the benefits of non-marginal reductions in pollution using information on defensive expenditures,” *Journal of Environmental Economics and Management*, 15, 111–127.

BENNEAR, L., K. JESSOE, AND S. OLMSTEAD (2009): “Sampling Out: Regulatory Avoidance and the Total Coliform Rule,” *Environmental science & technology*, 43, 5176–82.

BENNEAR, L. S. AND S. M. OLMSTEAD (2008): “The impacts of the “right to know”: Information disclosure and the violation of drinking water standards,” *Journal of Environmental Economics and Management*, 56, 117–130.

BERRY, S. T. AND G. COMPIANI (2020): “An Instrumental Variable Approach to Dynamic Models,” Working Paper 27756, National Bureau of Economic Research.

BLEVINS, J. R. (2016): “Sequential Monte Carlo Methods for Estimating Dynamic Microeconomic Models,” *Journal of Applied Econometrics*, 31, 773–804.

BLUNDELL, W., G. GOWRISANKARAN, AND A. LANGER (2020): “Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations,” *American Economic Review*, 110, 2558–85.

Brooks, L. and Z. Liscow (2023): "Infrastructure Costs," *American Economic Journal: Applied Economics*, 15, 1–30.

Christensen, P., D. Keiser, and G. Lade (2021): "Economic effects of environmental crises: Evidence from Flint, Michigan," *SSRN Electronic Journal*.

Collier, S. A., L. Deng, E. A. Adam, K. M. Benedict, E. M. Beshearse, A. J. Blackstock, B. B. Bruce, G. Derado, C. Edens, K. E. Fullerton, J. W. Gargano, A. L. Geissler, A. J. Hall, A. H. Havelaar, V. R. Hill, R. M. Hoekstra, S. C. Reddy, E. Scallan, E. K. Stokes, J. S. Yoder, and M. J. Beach (2021): "Estimate of burden and direct healthcare cost of infectious waterborne disease in the United States," *Emerg. Infect. Dis.*, 27, 140–149.

Congressional Budget Office (2018): "Public Spending On Transportation and Water Infrastructure, 1956-2017," .

Connault, B. (2016): "Hidden Rust Models," *Draft*.

Environmental Protection Agency (2002): "The Clean Water and Drinking Water Infrastructure Gap Analysis," .

——— (2008): "2006 Drinking Water Data Reliability Analysis and Action Plan For State Reported Public Water System Data in the EPA Safe Drinking Water Information System / Federal Version (SDWIS/FED)," .

——— (2009): "The Public Notification Rule: A Quick Reference Guide," .

Ericson, R. and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *The Review of Economic Studies*, 62, 53–82.

Graff Zivin, J., M. Neidell, and W. Schlenker (2011): "Water Quality Violations and Avoidance Behavior: Evidence from Bottled Water Consumption," *American Economic Review*, 101, 448–53.

Grooms, K. K. (2016): "Does Water Quality Improve When a Safe Drinking Water Act Violation Is Issued? A Study of the Effectiveness of the SDWA in California," *The B.E. Journal of Economic Analysis & Policy*, 16, 1–23.

Jardine, A. K., D. Lin, and D. Banjevic (2006): "A review on machinery diagnostics and prognostics implementing condition-based maintenance," *Mechanical Systems and Signal Processing*, 20, 1483–1510.

Jerch, R. L. (2022): "The Local Benefits of Federal Mandates: Evidence from the Clean Water Act," *Draft*.

Kalouptsidi, M., P. T. Scott, and E. Souza-Rodrigues (2021): "Linear IV regression estimators for structural dynamic discrete choice models," *Journal of Econometrics*, 222, 778–804.

KANG, K. AND B. S. SILVEIRA (2021): "Understanding Disparities in Punishment: Regulator Preferences and Expertise," *Journal of Political Economy*, 129, 2947–2992.

KEISER, D. A., B. MAZUMDER, D. MOLITOR, AND J. S. SHAPIRO (2023): "Water Works: Causes and Consequences of Safe Drinking Water in America," *Draft*.

KEIZER, M. C. O., S. D. P. FLAPPER, AND R. H. TEUNTER (2017): "Condition-based maintenance policies for systems with multiple dependent components: A review," *European Journal of Operational Research*, 261, 405–420.

KIM, M. J. AND V. MAKIS (2013): "Joint Optimization of Sampling and Control of Partially Observable Failing Systems," *Operations Research*, 61, 777–790.

KURT, M. AND J. P. KHAROUFEH (2010): "Optimally maintaining a Markovian deteriorating system with limited imperfect repairs," *European Journal of Operational Research*, 205, 368–380.

LEISTEN, M. AND N. VREUGDENHIL (2023): "Dynamic Regulation with Firm Linkages: Evidence from Texas," *Draft*.

LIM, C. S. H. AND A. YURUKOGLU (2018): "Dynamic Natural Monopoly Regulation: Time Inconsistency, Moral Hazard, and Political Environments," *Journal of Political Economy*, 126, 263–312.

MANSON, S., J. SCHROEDER, D. VAN RIPER, T. KUGLER, AND S. RUGGLES (2021): "National Historical Geographic Information System: Version 16.0 [dataset]," .

MARCUS, M. (2022): "Testing the Water: Drinking Water Quality, Public Notification, and Child Outcomes," *The Review of Economics and Statistics*, 104, 1289–1303.

NATIONAL RESEARCH COUNCIL (2006): *Drinking Water Distribution Systems: Assessing and Reducing Risks*, Washington, DC: National Academies Press.

NORETS, A. (2009): "Inference in Dynamic Discrete Choice Models with Serially Correlated Unobserved State Variables," *Econometrica*, 77, 1665–1682.

ORGILL-MEYER, J., M. JEULAND, J. ALBERT, AND N. CUTLER (2018): "Comparing Contingent Valuation and Averting Expenditure Estimates of the Costs of Irregular Water Supply," *Ecological Economics*, 146, 250–264.

POSENAU, K. (2022): "Debt and Water: Effects of Bondholder Protections on Public Goods," *Draft*.

REICH, G. (2018): "Divide and conquer: Recursive likelihood function integration for hidden Markov models with continuous latent variables," *Operations research*, 66, 1457–1470.

REYNOLDS, K. A., K. D. MENA, AND C. P. GERBA (2008): "Risk of waterborne illness via drinking water in the United States," *Rev. Environ. Contam. Toxicol.*, 192, 117–158.

RUST, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, 55, 999–1033.

SEARS, L. S., C.-Y. C. LIN LAWELL, G. TORRES, AND M. T. WALTER (2022): "Moment-based Markov Equilibrium Estimation of High-Dimension Dynamic Games: An Application to Groundwater Management in California," *Draft*.

SI, X.-S., W. WANG, C.-H. HU, AND D.-H. ZHOU (2011): "Remaining useful life estimation–a review on the statistical data driven approaches," *European journal of operational research*, 213, 1–14.

TIMMINS, C. (2002): "Measuring the Dynamic Efficiency Costs of Regulators' Preferences: Municipal Water Utilities in the Arid West," *Econometrica*, 70, 603–629.

WRENN, D. H., H. A. KLAIBER, AND E. C. JAENICKE (2016): "Unconventional shale gas development, risk perceptions, and averting behavior: Evidence from bottled water purchases," *Journal of the Association of Environmental and Resource Economists*, 3, 779–817.

WU, P.-I. AND C.-L. HUANG (2001): "Actual averting expenditure versus stated willingness to pay," *Applied Economics*, 33, 277–283.

Appendix

A Project Classification

Prior to the analysis of the relationship between infrastructure projects and violations, I classify projects as *reactive* or non-reactive. I consider all non-reactive projects to be *proactive*. The primary purpose of this exercise is to isolate the projects that are pursued in response to a system issue relevant to the provisioning of safe drinking water. Project reports include detailed project descriptions and "need for project" justifications that I use to classify projects. I consider a project to be reactive if either the description or the project need contains language alluding to a public health emergency or a system failure. For example, any excerpt that mentions the system facing a public health problem, needing to return to compliance, or a major infrastructure issue that affects whether there are contaminants in the water (e.g., pipe deterioration, extensive breaks, or overstressing) are all considered to be reactive. See Table A.1 for examples of project excerpts and their classifications.

I employ natural language processing (NLP) tools to categorize the entire population of projects. To this end, I first manually label 250 project descriptions and needs, assigning each passage as reactive or not.[37] Approximately 26% of the initial 250 assignments are reactive. I then use the manual classifications to train an NLP model. In the training process, the program constructs a statistical model to predict the reactive categorization by adjusting a series of weights placed on a set of unobserved factors. The program determines the values of the model weights by splitting the manually labeled data into different groups and verifying the calculated weights can predict the provided classifications. The NLP model has an AUC value of 0.881, indicating that the classification made correct predictions in approximately 88.1% of the evaluation cases in my training sample.

I then use the NLP model on the unlabeled dataset to assign projects as reactive or proactive. At this time, I also exclude extension projects, projects that add new customers to the water system.[38] Approximately 30% of all projects are extensions and are not included in further analysis. I assign projects as reactive based on the model's classification of the text in both the project description and project need. If both of these passages have a reactive probability classification at 25% or below, the project is classified as proactive. If either the project description or project need has a reactive probability classification at 75% or above, the project is classified as reactive. Figure A.1 depicts

[37] When possible, I validated my manual classification with the occasionally-completed *DW Specific Impacts* checklist contained within the project reports. The checklist indicated if the project was in response to an emergency, required for compliance, or for anticipated future requirements.

[38] I identify extension projects by nonzero values in the "Total New Households" field included in the project reports. These extension projects are a direct reflection of a 1999 executive order to provide the "best available water and sewer service to every Kentuckian by 2020." As the motivation for these projects is not to maintain existing infrastructure, I exclude them from my analysis.

Table A.1: Project Classification Examples

Project	Project Text	Classification
WX21103050 ($102,670.20)	"The city needs to replace sections of old cast iron (ci) and asbestos cement (AC) water line; as well as, inoperable gate valves throughout its water system. A majority of the water system is 60-80 years old ci and is corroded and constricts flows. The AC line breaks easily and causes undue concern to the citizens due to the asbestos content in the pipe."	Reactive
WX21135016 ($119,513.26)	"Utilities were damaged due to flooding. Repairs are necessary to maintain service."	Reactive
WX21125552 ($8,673,494.55)	"The project will construct additional filters, controls filter building, drying beds, polymer feed equipment and associated piping and appurtenances to upgrade the plant operation to 3.0 MGD."	Proactive
WX21173050 ($49,357)	"The project will increase water pressure, improve customer service, water quality, and water delivery."	Proactive

Notes: Project costs are adjusted to real 2012 dollars. Classification for project WX21135016 is manually assigned, all others assigned by the NLP model.

the histogram of reactive probabilities for project descriptions and needs. The high concentration of probabilities around 0 and 1 indicate there were few cases in which the NLP model had difficulty classifying the text. I hand label the remaining projects that have unclear reactive assignments based on the model classification. After removing extensions and projects with missing data, 2,169 projects remain. The remaining project population is about 30% reactive, which is slightly larger than the fraction in the original hand labeled set.

B Recursive Likelihood Algorithm Details

The likelihood function for an individual water system can be represented as:

$$\begin{aligned} L_T(\boldsymbol{\theta}) &= \int \dots \int p_q(q_0;\boldsymbol{\theta}) \\ &\times \quad P_{iy_r k_p k_r y_v q}\Big(\{i_t, y_{rt}, k_{pt}, k_{rt}, y_{vt}, q_t\}_{t=1}^{T} | \{i_0, y_{r0}, k_{p0}, k_{r0}, y_{v0}, q_0\}; \boldsymbol{\theta}\Big) dq_0 \dots dq_T \end{aligned} \tag{B.1}$$

All transitions except for the progression of q, are independent, and the progression of q is determined by equation (2) as outlined in Section 4. As a result, the joint probability can be rewritten as follows.

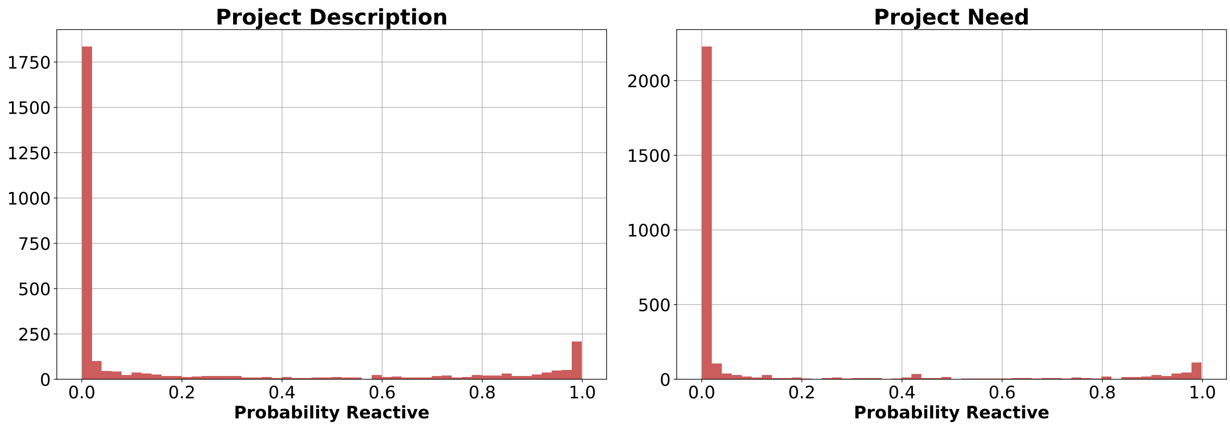

Figure A.1: Histogram of Reactive Probabilities Predicted by NLP Model
Notes: Data are from WRIS and collected via web scraping. Reactive probabilities are assigned using the NLP model.

$$\begin{aligned} &P_{iy_rk_pk_ry_vq}\left(\{i_t, y_{rt}, k_{pt}, k_{rt}, y_{vt}, q_t\}_{t=1}^{T}|\boldsymbol{\theta}\right) = \\ &\prod_{t=1}^{T} p_{iy_rk_pk_ry_vq}(i_t, y_{rt}, k_{pt}, k_{rt}, y_{vt}, q_t | i_{t-1}, y_{rt-1}, k_{pt-1}, k_{rt-1}, y_{vt-1}, q_{t-1}; \boldsymbol{\theta}) \end{aligned} \tag{B.2}$$

Note that $p_{k_p}(k_{pt};\boldsymbol{\theta})$ and $p_{k_r}(k_{rt};\boldsymbol{\theta})$ are independent of q and are therefore pre-estimated and omitted from future representations of the likelihood function.

Incorporating the assumptions into the model, and applying Fubini's Theorem, the likelihood becomes:

$$\begin{aligned} L_T(\boldsymbol{\theta}) &= \int ... \int \prod_{t=1}^{T} \Big(p_{i|q}(i_t|q_t;\boldsymbol{\theta}) p_{y_r|q}(y_{rt}|q_t, i_t;\boldsymbol{\theta}) p_{y_v|q}(y_{vt}|q_t;\boldsymbol{\theta}) \\ &\times \quad p_q(q_t|i_{t-1}, y_{rt-1}, k_{pt-1}, k_{rt-1}, q_{t-1};\boldsymbol{\theta}) \Big) \\ &\times \quad p_{i|q}(i_0|q_0;\boldsymbol{\theta}) p_{y_r|q}(y_{r0}|q_0, i_0;\boldsymbol{\theta}) p_{y_v|q}(y_{v0}|q_0;\boldsymbol{\theta}) p_q(q_0;\boldsymbol{\theta}) dq_T dq_{T-1} ... dq_0 \end{aligned} \tag{B.3}$$

Rewriting the likelihood function:

$$
\begin{aligned}
L_T(\boldsymbol{\theta}) &= \int ... \int \Bigg[\prod_{t=1}^{T-1} \Bigg(p_{i|q}(i_t|q_t;\boldsymbol{\theta}) p_{y_r|q}(y_{rt}|q_t, i_t;\boldsymbol{\theta}) p_{y_v|q}(y_{vt}|q_t;\boldsymbol{\theta}) \\
&\times \; p_q(q_t|i_{t-1}, y_{rt-1}, k_{pt-1}, k_{rt-1}, q_{t-1};\boldsymbol{\theta}) \Bigg) \\
&\times \; p_{i|q}(i_0|q_0;\boldsymbol{\theta}) p_{y_r|q}(y_{r0}|q_0, i_0;\boldsymbol{\theta}) p_{y_v|q}(y_{v0}|q_0;\boldsymbol{\theta}) p_q(q_0;\boldsymbol{\theta}) \Bigg] \\
&\times \; \Bigg(\int p_{i|q}(i_T|q_T;\boldsymbol{\theta}) p_{y_r|q}(y_{rT}|q_T, i_T;\boldsymbol{\theta}) p_{y_v|q}(y_{vT}|q_T;\boldsymbol{\theta}) \\
&\times \; p_q(q_T|i_{T-1}, y_{rT-1}, k_{pT-1}, k_{rT-1}, q_{T-1};\boldsymbol{\theta}) dq_T \Bigg) dq_{T-1}...dq_0
\end{aligned} \tag{B.4}
$$

Define the following at time $t = T$:

$$
\begin{aligned}
&E\Big[p_{i|q}(i_T|\tilde{q}_T;\boldsymbol{\theta}) p_{y_r|q}(y_{rT}|\tilde{q}_T, i_T;\boldsymbol{\theta}) p_{y_v|q}(y_{vT}|\tilde{q}_T;\boldsymbol{\theta}) \Big| i_{T-1}, y_{rT-1}, k_{pT-1}, k_{rT-1}, q_{T-1};\boldsymbol{\theta} \Big] = \\
&\quad \int p_{i|q}(i_T|q_T;\boldsymbol{\theta}) p_{y_r|q}(y_{rT}|q_T, i_T;\boldsymbol{\theta}) p_{y_v|q}(y_{vT}|q_T;\boldsymbol{\theta}) p_q(q_T|i_{T-1}, y_{rT-1}, k_{pT-1}, k_{rT-1}, q_{T-1};\boldsymbol{\theta}) dq_T
\end{aligned} \tag{B.5}
$$

and at time $t = T - 1$:

$$
\begin{aligned}
&E\Bigg[\prod_{t=T-1}^{T} p_{i|q}(i_T|\tilde{q}_T;\boldsymbol{\theta}) p_{y_r|q}(y_{rT}|\tilde{q}_T, i_T;\boldsymbol{\theta}) p_{y_v|q}(y_{vT}|\tilde{q}_T;\boldsymbol{\theta}) \Big| i_{T-2}, y_{rT-2}, k_{pT-2}, k_{rT-2}, q_{T-2};\boldsymbol{\theta} \Bigg] = \\
&\int E\big[p_{i|q}(i_T|\tilde{q}_T;\boldsymbol{\theta}) p_{y_r|q}(y_{rT}|\tilde{q}_T, i_T;\boldsymbol{\theta}) p_{y_v|q}(y_{vT}|\tilde{q}_T;\boldsymbol{\theta}) \big| i_{T-1}, y_{rT-1}, k_{pT-1}, k_{rT-1}, q_{T-1};\boldsymbol{\theta} \big] \\
&\qquad \times \Big(p_{i|q}(i_{T-1}|q_{T-1};\boldsymbol{\theta}) p_{y_r|q}(y_{rT-1}|q_{T-1}, i_{T-1};\boldsymbol{\theta}) p_{y_v|q}(y_{vT-1}|q_{T-1};\boldsymbol{\theta}) \\
&\qquad \times p_q(q_{T-1}|i_{T-2}, y_{rT-2}, k_{pT-2}, k_{rT-2}, q_{T-2};\boldsymbol{\theta}) dq_{T-1} \Big)
\end{aligned} \tag{B.6}
$$

Using this notation and working backward from $t = T$, the likelihood can be represented as:

$$
\begin{aligned}
L_T(\boldsymbol{\theta}) &= \int E\Bigg[\prod_{t=1}^{T} p_{i|q}(i_t|\tilde{q}_t;\boldsymbol{\theta}) p_{y_r|q}(y_{rt}|\tilde{q}_t, i_t;\boldsymbol{\theta}) p_{y_v|q}(y_{vt}|\tilde{q}_t;\boldsymbol{\theta}) \Big| q_0;\boldsymbol{\theta} \Bigg] \\
&\times \; p_{i|q}(i_0|q_0;\boldsymbol{\theta}) p_{y_r|q}(y_{r0}|q_0, i_0;\boldsymbol{\theta}) p_{y_v|q}(y_{v0}|q_1;\boldsymbol{\theta}) p_q(q_0;\boldsymbol{\theta}) dq_0
\end{aligned} \tag{B.7}
$$

Table C.1: Simulations Demonstrating Non-Identification of μ_q

Parameter		True Values	Baseline	High Mean	Low Mean
Delay Consequences					
Consumer Cost Weight	λ	10	9.707	8.061	13.602
			(1.000)	(0.301)	(0.999)
Reactive Project Threshold	q_r^*	47	47.156	65.307	26.843
			(1.045)	(0.091)	(1.720)
1Q Violation Threshold	q_{v0}^*	45	44.491	62.440	24.155
			(0.665)	(0.080)	(0.760)
2Q Violation Threshold	q_{v1}^*	44	43.185	61.440	23.147
			(1.000)	(0.078)	(0.665)
3Q Violation Threshold	q_{v2}^*	43	41.843	60.440	22.117
			(0.937)	(0.075)	(0.622)
4Q Violation Threshold	q_{v3}^*	42	40.557	59.441	21.061
			(0.916)	(0.076)	(0.602)
Project Effect on Quality					
Proactive Investment	α^p	0.1	0.120	0.158	0.097
			(0.025)	(0.006)	(0.026)
Reactive Investment	α^r	0.05	0.054	0.081	0.026
			(0.006)	(0.003)	(0.005)
Initial Quality Distribution					
Initial Quality Variance	σ_q^2	1.5	1.393	2.296	1.740
			(1.002)	(0.157)	(1.333)
Fixed Parameters					
Initial Quality Average	μ_q	50	50	70	30
Deterioration Rate	α^q	0.99	—	—	—
Neg. Log Likelihood					
ℓ		17177.967	16967.529	16917.805	17184.885

C Identification Simulations

In this section, I present the results from simulations that examine the identification of the mean of the initial distribution of quality, μ_q. I use the assumed data generating process to simulate a sample dataset for 353 water systems over 50 years with known parameters of interest. I then ran the RLI algorithm on the simulated data to recover estimates for the parameters. I first ran the estimation process holding μ_q fixed at the true value used for simulation to determine if the program is capable of recovering the correct values for the other parameters. I then ran the algorithm on the data again but held μ_q fixed at incorrect values. Table C.1 contains a summary of my results.

The "Baseline" column depicts the results of the estimation when holding μ_q fixed at the same value used to generate the data. Under these conditions, the RLI algorithm is able to recover the remaining parameters. The algorithm slightly underestimates the threshold parameters but overall, the estimates are close to the true values. In the other specifications, I hold μ_q fixed above the true value, and below the true value to determine the effect on the parameter estimates. Based on these tests, I conclude that the mean of the q_0 distribution determines the relative levels of many of the other parameters in the model. When holding $\mu_q = 70$, the thresholds for violation duration and reactive projects increase to match the observations of the number of periods that pass before a system spends time in violation. Similarly, when holding $\mu_q = 30$, the thresholds for violation duration and reactive projects decrease. Additionally, the estimates for the effects of infrastructure projects on quality also follow the change in the initial quality level and the consumer cost weight moves in the opposite direction. Based on this evidence, I conclude that μ_q is not identifiable, and I hold this value fixed in the estimation of the model.

D Robustness Checks

D.1 System Manager Foresight

In the model presented in Section 4, I assume that managers are forward-looking. To justify investment decisions that have limited immediate benefits, system managers must value the future state of the system. I also assume that system managers make annual decisions considering the infinite horizon of possible future states of their quality and that they discount the future at a rate of $\beta = 0.95$. In this section, I examine the results of alternative assumptions to assess the possibility that system managers make decisions under different conditions. I focus my evaluation on the changes in the effect of each type of project on quality, (α^p, α^r), and the weight that system managers place on consumer exposure costs, λ. The thresholds for violation and reactive projects are not identified by system manager decisions and therefore remain relatively unchanged in alternative estimations.

First, I analyze results from models in which system managers make one-shot decisions based on projections of the state of their infrastructure at a fixed number of years in the future. The main difference between these models and the baseline model is in the calculation of the probability that a system manager invests in an infrastructure project. I consider specifications where system managers assess their infrastructure state one, five, ten, or fifteen years in the future. The results of these estimations are presented in Table D.1. The one-shot model with a one-year time horizon is the worst match to the data. This specification has the highest objective value, implying the worst fit, and overestimates both the effectiveness of projects and the weight placed on consumer exposure costs. With the exception of the ten-year time horizon, the parameter estimates and the negative log likelihood function approach values that are close to the baseline specification as the time horizon for the one-shot specifications increases. The ten-year time horizon appears to perform slightly better than the infinite horizon model although the differences in objective functions are relatively minor.

Second, I alter the rate at which system managers discount the future. The discount factor

Table D.1: Time Horizon Comparison

	Baseline	1 Year	5 Year	10 Year	15 Year
α^p	0.061	0.397	0.113	0.043	0.094
	(0.006)	(0.034)	(0.008)	(0.010)	(0.041)
α^r	0.038	0.038	0.028	0.011	0.038
	(0.005)	(0.016)	(0.003)	(0.001)	(0.011)
λ	9.522	56.808	29.639	41.485	12.411
	(1.014)	(1.049)	(0.977)	(1.000)	(1.090)
Neg. Log Likelihood					
ℓ	5734.038	5935.766	5760.516	5725.053	5735.433

Table D.2: Discount Factor Comparison

	Baseline	$\beta = 0.8$	$\beta = 0.9$	$\beta = 0.99$
α^p	0.061	0.139	0.100	0.064
	(0.006)	(0.018)	(0.019)	(0.017)
α^r	0.038	0.034	0.039	0.063
	(0.005)	(0.004)	(0.006)	(0.017)
λ_0	9.522	24.979	12.901	4.308
	(1.014)	(1.195)	(1.543)	(1.025)
Neg. Log Likelihood				
ℓ	5734.038	5762.545	5736.300	5736.047

describes the relationship between current and future events. In the baseline specification, I hold this parameter fixed at 0.95, meaning that to a system manager a $1 expenditure tomorrow is only worth $0.95 today. To compare alternative discount rates, I consider specifications where the discount factor is held fixed at 0.8, 0.9, and 0.99. The result of these estimations are presented in Table D.2. The baseline specification fits the data better than the alternative specifications, although the value of the negative log likelihood function is similar in all four models. In the specification with $\beta = 0.8$, the estimate for the weight system managers place on consumer costs increases to compensate for the lower discount rate. Similarly, in the specification with an increased discount rate of $\beta = 0.99$, estimates for λ are lower as system managers in this model place a higher weight on the future. Across the model specifications I consider, I find little conclusive evidence that an alternative model for system manager decisions provides substantial improvement over the baseline infinite horizon model. The baseline model has the most believable estimates for the weight placed on consumer exposure costs, and this specification precisely estimates the parameter values.

E Additional Figures

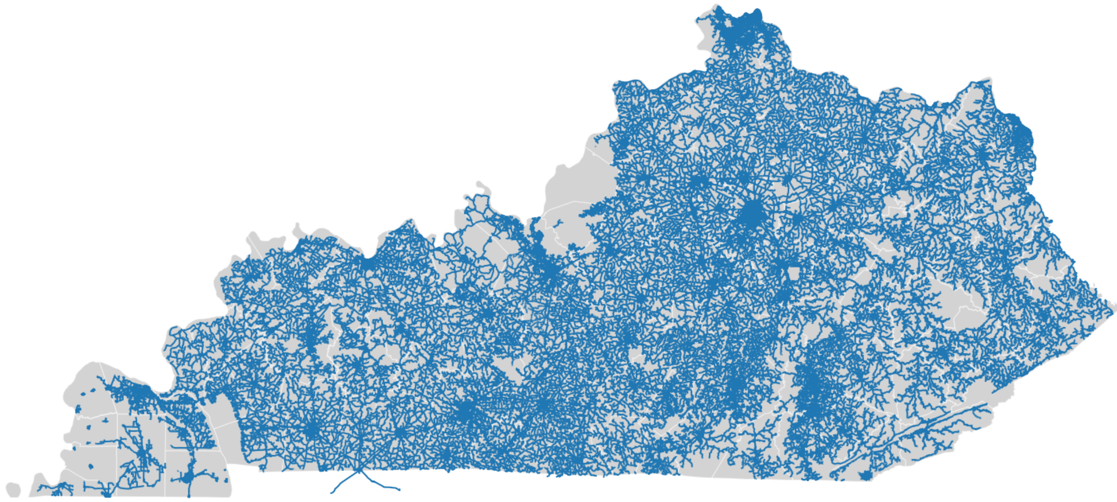

Figure E.1: Map of Water Lines Comprising Kentucky's Public Water Systems
Notes: Data compiled from geojson files obtained from the geographic information system (GIS) portion of WRIS, collected in July 2021. Water lines for the Louisville Water Company, which provides service to Louisville, are unavailable.

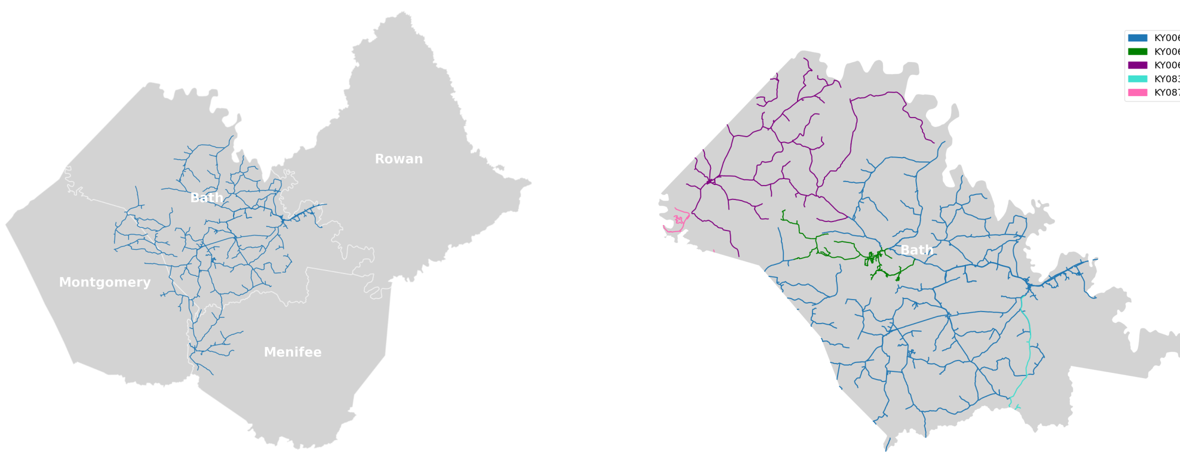

(a) Counties Served by KY0060022 (b) Bath County Water Systems

Figure E.2: County Community Water System Overlays

Notes: Data are compiled from WRIS geojson files. Frequently water system geographies do not follow county lines. As a result, multiple systems often serve a single county and a system often serves populations across multiple counties.

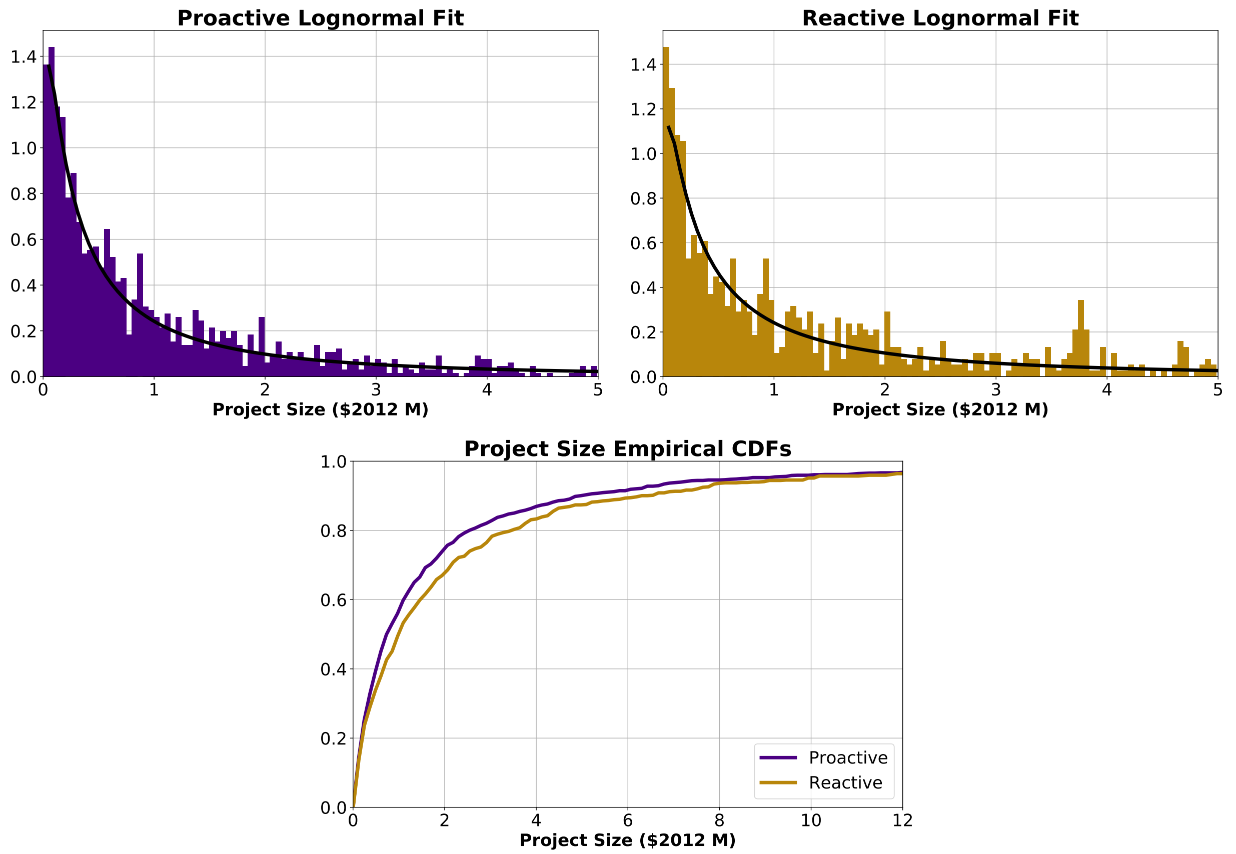

Figure E.3: Lognormal Fits for Project Investments

Notes: Data are from WRIS and collected via web scraping. Project costs are adjusted to millions of real 2012 dollars. The top two panels plot the histograms of proactive and reactive projects along with the estimated lognormal fit. The bottom panel plots the empirical cumulative distributions for both project types.

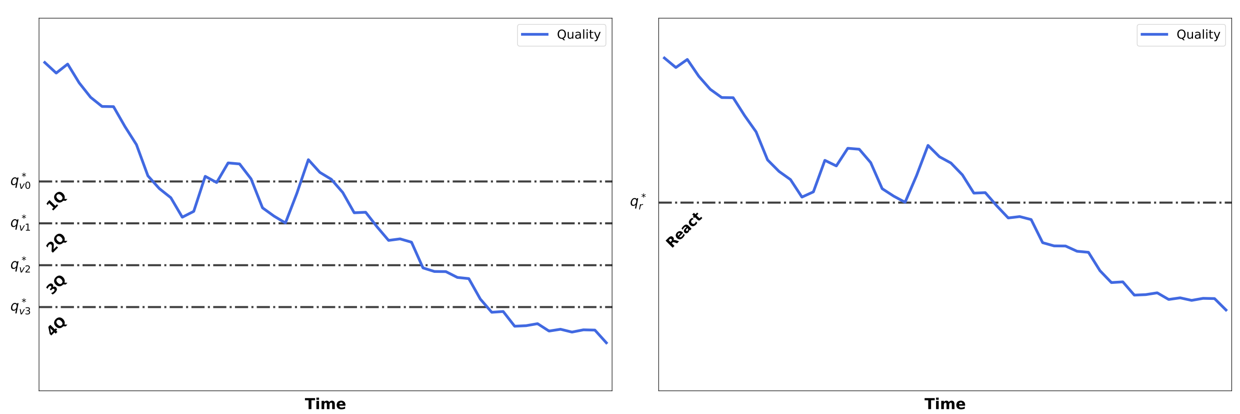

Figure E.4: Low Quality Consequences

Notes: The left panel portrays the relationship between quality decline and violation thresholds. The right panel portrays the relationship between quality decline and the reactive threshold. All representations are simulated based on the structure of the model.

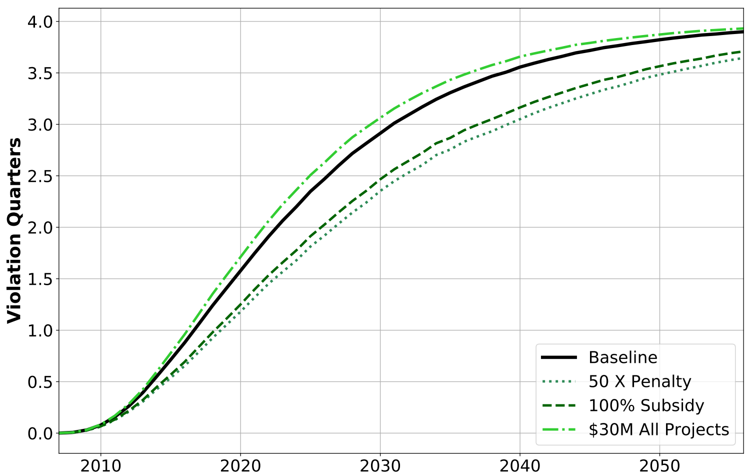

Figure E.5: Counterfactual Expected Violation Quarters

Notes: The figure depicts the simulated predicted average number of quarters spent in violation across all public water systems over 2007-2057.

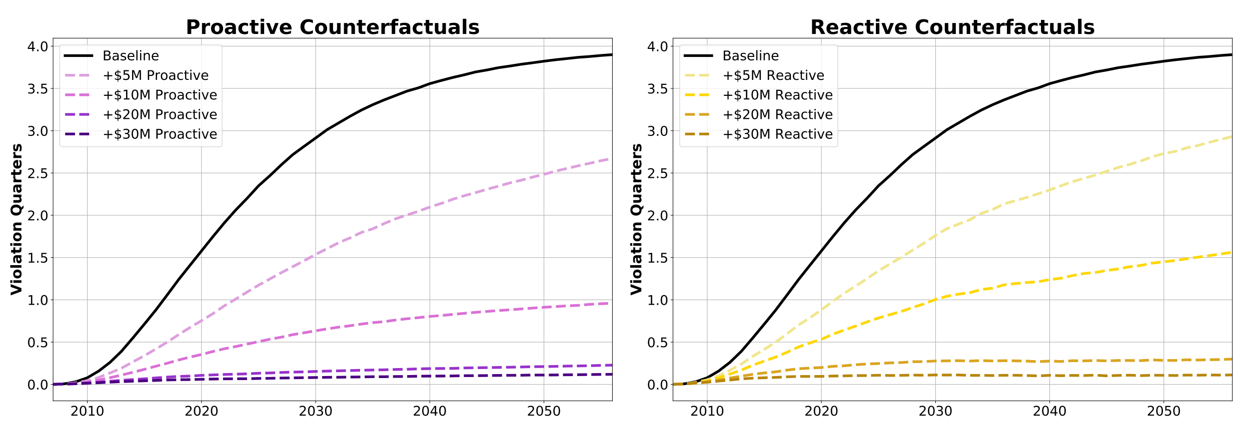

Figure E.6: Counterfactual Expected Violation Quarters, Increased Investment

Notes: The left panel displays the expected violation quarters from 2007-2057 for different proactive investment levels, discounted to be free to system managers. The right panel displays the same information for reactive investments. The black solid line depicts the expected violation quarters in the baseline model.