Does Information Disclosure Reduce Drinking Water Violations in the United States?

Jonathan Baker, Lori Bennear, and Sheila Olmstead[1]

August 29, 2022

Final version published 2023 in the
Journal of the Association of Environmental and Resource Economists

Abstract

The 1996 Safe Drinking Water Act Amendments required community water systems to disclose violations of drinking water standards to their customers in annual water quality reports. We explore the impact of three methods of disclosure on health-based drinking water quality violations using a matching and differences-in-differences framework with a national dataset of drinking water quality violations from 1990-2001. We find that this information disclosure requirement reduced drinking water violations significantly, and that the primary effect of disclosure on violations persists for at least four years after policy implementation. We find no evidence, however, that water systems trade these potentially more salient violation reductions for potentially less salient reductions in violations of other standards, nor do we find any evidence that water systems responded differentially to disclosure based on the demographic or political characteristics of their customers.

[1] Baker (jonathan.baker@analysisgroup.com) is a Manager at Analysis Group. Bennear (lori.bennear@duke.edu) is the Juli Plant Granger Associate Professor of Energy Economics and Policy and the Executive Vice Dean at Duke University's Nicholas School of the Environment. Olmstead (sheila.olmstead@austin.utexas.edu) is a Professor of Public Affairs at the University of Texas at Austin's LBJ School of Public Affairs, a University Fellow at Resources for the Future, and a Senior Fellow at the Property and Environment Research Center. The authors thank EPA staff for providing the water systems and water violations data, as well as many helpful discussions about the data. We also gratefully acknowledge Simo Goshev of Harvard's Institute for Quantitative Social Science for his advice and guidance on aspects of our empirical strategy, and seminar participants at the Harvard Environmental Economics Lunch, Harvard Public Policy API-902, and the University of Delaware Economics Department and School of Marine Science and Policy. Alberto Abadie, Matt Baum, Todd Gerarden, Nick Hagerty, Robert Stavins, and Richard Zeckhauser also provided helpful input.

1. Introduction

The recent crisis in Flint, Michigan concerning toxic levels of lead in the city's drinking water thrust water quality into the national spotlight. Since the passage of the Safe Drinking Water Act (SDWA) in 1974, the U.S. Environmental Protection Agency has endeavored to limit the occurrence of Flint-like crises by regulating drinking water quality across the country's tens of thousands of public water systems (Tiemann 2017). In 1996, Congress substantially amended the SDWA. One new requirement mandated that community water systems—those serving a non-transient population more than 6 months per year— submit annual water quality reports to their customers. The move in part reflected a trend among regulators to achieve policy goals through information disclosure, especially in contexts where direct, centralized regulatory control poses challenges (Fung et al. 2007). While all community water systems must generate reports, the method of disclosure required by the 1996 Amendments depends upon the water system size. Systems serving up to 500 customers must make a report available upon request; systems serving between 501 and 9,999 customers must publish a report in a local newspaper; systems serving at least 10,000 customers must mail the report; and systems serving at least 100,000 customers must mail a report and also make the report available on the Internet. The first reporting year was 1998.

We estimate the change in health-based water quality violations in the United States due to the annual disclosure of water quality reports required by the 1996 SDWA Amendments. Using data we collected on all U.S. water systems and all water quality violations (1990-2001) from the U.S. Environmental Protection Agency (EPA) through a Freedom of Information Act request, we build a panel of community water systems active from 1990 through 2001, matching the number of health-based water quality violations with the violating water system in each year. We construct three panels for each information disclosure threshold where, in each panel, we match treated with untreated water systems on system characteristics and pre-treatment average annual health-based water quality violations. To estimate the impact of the publishing, mailing, and online posting disclosure requirements, we apply a series of differences-in-differences models that exploit the timing of the information disclosure policy and the discontinuity in the required method of disclosure. We find that the fact rather than the method of information disclosure drives water systems' response, with disclosure overall reducing water quality violations by around 33 percent relative to pre-disclosure policy levels. This basic result is robust to many different model specifications. While some models suggest that the direct mailing and online posting requirements may also have had some incremental effect on violations, these results are less robust than those for the baseline disclosure requirement.

The effects of information disclosure have been studied in a wide range of contexts, including finance, health, and education (Dranove and Jin 2010), energy (Allcott and Sweeney 2016), and the environment (Mastromonaco 2015). The paper most closely related to ours is Bennear and Olmstead (2008), which estimates that the 1996 SDWA requirement that water suppliers mail violation information to served households reduced total violations in Massachusetts by 30 – 44 percent, with somewhat stronger effects for health-based violations. However, this prior paper uses data from only one U.S. state. Thus, one contribution of our work is a more comprehensive picture of the impacts of the SDWA's information disclosure requirements – a federal regulation – on violations nationwide. In addition, in the prior Massachusetts-only analysis, the number of water suppliers in each of the reporting requirement bins was too small to estimate separately the effects of each form of disclosure, and violations for contaminants other than bacteria were insufficient to estimate the impacts of disclosure on individual contaminants. We examine both of these questions in our analysis.

To examine impacts of disclosure on violations of specific contaminant types, we focus on violations pertaining to microbial contaminants (such as cryptosporidium or *E. coli*) and contaminants created as byproducts of the disinfection process (primarily trihalomethanes). These two violation types are related in a technical sense because the chemicals used to disinfect drinking water (to kill microbes) interact with other substances in the treatment and distribution system to create trihalomethanes and other disinfection byproducts (DBPs). Pathogenic microbes cause acute illness, and DBPs are potential carcinogens – both are regulated under the SDWA. Thus, there is a direct tradeoff between disinfecting drinking water supplies so as to reduce the probability that end-users are exposed to bacterial, viral and other pathogens, and the degree of exposure to carcinogens resulting from the disinfection process. Prior work suggests that microbial violations may be more salient than DBP or other violations to water users (Adamowicz et al. 2011, Graff Zivin et al. 2011). Adamowicz et al. (2011) find that individuals have higher willingness to pay to reduce microbial contaminants, which cause acute gastrointestinal illness, compared to known carcinogens with latent effects. As a consequence, water systems may be more eager to reduce microbial contaminants than other contaminants that increase the risk of cancer or contribute to other long-term deleterious health conditions. Water suppliers may therefore respond to disclosure by trading one type of violation for another. Using a seemingly unrelated regression framework, we find no evidence that the disclosure requirement increases DBP violations when it decreases microbial violations nationwide. Results are similar but less precise for the mailing and online disclosure requirements.

Another important contribution is our analysis of persistence. Little empirical work describes the persistence of impacts from information disclosure policies (Dranove and Jin 2010, Balasubramanya et al. 2014). We address this gap in the literature, finding that violation reductions caused by disclosure under the SDWA are stable and persist for at least the first four years of the policy. This finding is related to recent findings on the persistence of consumers' responses to behavioral interventions (especially social comparisons) to induce energy and water conservation (Allcott and Rogers 2014, Ferraro et al. 2011), as well as regulated entities' responses to such interventions (Earnhart and Ferraro 2021). However, the context of information disclosure about violations of an environmental regulation differs significantly from social comparison policies, which affect voluntary behavior. The persistent impacts we estimate suggest that water utilities may have made permanent changes to their compliance strategies in response to the disclosure requirement.

We also contribute to the limited literature on heterogeneity in responsiveness to information disclosure. Delmas et al. (2010) show that investor-owned utilities shift away from fossil fuels to cleaner fuels when required to disclose their fuel mixture, and that this effect is more pronounced for utilities serving primarily residential customers. Powers et al. (2011) find that large pulp and paper manufacturers in India reduce pollution in response to an environmental quality rating program, especially in higher income areas. Shimshack et al. (2007) find that the effects of mercury warnings on fish consumption vary with consumers' income and newspaper readership.[2] More broadly, the environmental justice literature examines links between pollution exposure, income, race, and other demographic characteristics. We test for heterogeneous responses to

[2] While not related to information disclosure, other work has found that community characteristics influence compliance with environmental regulations by publicly-owned wastewater treatment facilities (Earnhart 2004).

information disclosure as a function of the demographic composition of a water system's service population, and we find no evidence of a differential responses to disclosure along these lines. While a contrast to some prior work on information disclosure, our result is consistent with other papers in the economics literature suggesting that some environmental policies have neutral or even positive environmental justice outcomes (Fowlie et al. 2012, Mansur and Sheriff 2021). While minority and low-income communities do appear to bear a disproportionate burden from drinking water violations (Allaire et al. 2018), our work does not suggest that the requirement to disclose information about violations to consumers on an annual basis has exacerbated this inequity. Finally, we also test for differential responsiveness to disclosure based on the political affiliation of water systems' service population, finding no significant impact.

Several papers in the literature suggest possible mechanisms through which information disclosure could affect behavior (Sunstein 1999, Reinhart 2000, Fung and O'Rourke 2000, Karkkainen 2001). If information about water systems' environmental performance is known and consumers have preferences over environmental performance, systems can face market pressure to improve performance. This is less likely in the SDWA context because one-half of systems in our sample are publicly owned, most of the private systems are not-for-profit organizations, and essentially all of them are monopolies within their service areas. There is evidence in the literature (Graff Zivin et al. 2011) that households may substitute from tap water to bottled water in response to bacterial and other violations under the SDWA. However, drinking water represents a tiny fraction of the water sold by these systems, suggesting that the market mechanism is not a likely driver of the impacts of disclosure under the SDWA.[3] The fact that the disclosure effects we estimate are insensitive to household income also suggests that market pressure is an unlikely mechanism in this case.

A second possible mechanism is political; the CCR informs water users of a risk about which they were previously unaware (or under-aware), and those users exert pressure on the water

[3] According to the EPA, about 70 percent of household water use occurs indoors (presumably outdoor water is not used for drinking). According to the 2016 Residential End-Uses of Water Study (Water Research Foundation 2016), about 20 percent of indoor water use comes from faucets. According to that same study, one-third of faucet use is hot water, which is presumably not for drinking. This leaves about 9.2 percent (0.7 * 0.2 * 0.67) of total household water use as possible drinking water. However, this 9.2 percent also comprises water used for cooking, hand-washing, teeth-brushing, and other faucet uses. Conservatively, even if households used one-third of this (about 3 percent) of total water consumption for drinking, and substituted completely to bottled water, this would have a very small impact on water systems.

system to improve environmental performance. This is a possible mechanism for the reductions in violations seen in the context of the SDWA (Bennear and Olmstead 2008). However, our null findings with respect to service population demographics and political characteristics makes this unlikely.

The act of measuring and reporting data on environmental performance may itself provide new information to the water system, leading to internal changes that improve environmental performance. Risk-reducing process improvements are certainly a possible mechanism for the violation reductions we estimate, though without data on the specific processes followed by individual systems, we cannot test for this. A mechanism of learning and process improvement is perhaps most consistent with our results indicating that service population demographics and political affiliation do not impact the effects of disclosure on water system violations. It is also consistent with our results suggesting that reductions in violations persist over time.

One concern if process change is the mechanism for the results we estimate would be that water systems are responding strategically, reducing salient violations without necessarily improving public health. We test explicitly for one such possibility: that systems increase disinfection, raising the risk of violating DBP standards in order to reduce the risk of bacterial violations. We find no evidence for this behavior. Another possibility is that systems sample strategically to avoid the detection of a violation – a behavior observed in prior work (Bennear et al. 2009). However, testing for this behavior would require detailed individual paper records for each of the more than 46,000 systems in our sample, a data collection process that is beyond our scope.[4] Thus, while we view learning and process improvement in reducing the frequency of violations as perhaps the most likely driver of our results, understanding how information disclosure affects regulatory compliance in this context and others remains an important area for future research.

We organize the remainder of the paper as follows. In section 2 we discuss our policy context, the 1996 SDWA amendments. We describe our data in section 3 and our empirical models in section 4. We present and discuss results in section 5. Section 6 concludes the paper.

[4] Bennear et al. (2009) spent one year compiling the necessary data from paper files in the regional offices of the state DEP in Massachusetts. It is not feasible to acquire these data for every U.S. water system.

2. Policy Context: 1996 Safe Drinking Water Act Amendments

Since the passage of the SDWA in 1974, the EPA has regulated the quality of drinking water delivered by U.S. water utilities. Initially, the SDWA regulated 22 contaminants including lead, arsenic, coliform bacteria, and mercury. New contaminants, such as radioactive isotopes, copper, and cryptosporidium, have been added as the Act has been amended over time. In 2021, the EPA regulated more than 90 contaminants under the SDWA.[5]

In 1996 Congress undertook a major overhaul of the SDWA. Among other requirements, the 1996 Amendments instituted a rule stipulating that all community water systems submit annual water quality reports to their customers. Referred to as Consumer Confidence Reports (CCRs), these annual water quality report cards are intended to "Improve public health protection by providing educational material to allow consumers to make educated decisions regarding any potential health risks pertaining to the quality, treatment, and management of their drinking water supply."[6] In addition to including information about the water system and contaminants, the water quality report must make note of any water quality violations that occurred during the reporting period, and explain any potential health implications.

The required method of disseminating the annual water quality report depends upon the service population of the water system. Water systems serving few customers face the least stringent requirements; below a service population of 501, water systems can provide a notice stating the CCR is available upon request. Between a service population of 501 and 9,999 customers, water systems must make the report available in a public forum, for example, publishing their CCR in a local newspaper. Water systems must mail their report if they serve 10,000 or more customers, and systems serving 100,000 or more persons must also publicly post their current year's report on the Internet. We refer to these three information disclosure requirements as the publishing, mailing, and online posting disclosure requirements, respectively.[7]

[5] A timeline of regulated contaminants can be found at: https://www.epa.gov/sites/production/files/2015-10/documents/dw_regulation_timeline.pdf.

[6] Requirements of the CCR in this paragraph are summarized from: EPA Quick Reference Guide No. 816-F-09-009. http://nepis.epa.gov/Exe/ZyPDF.cgi?Dockey=P100529A.txt.

[7] In 2013, EPA began allowing systems in the "mailing" group to post their CCRs online as a substitute for mailing, so long as a notice about the CCR's availability and online location is mailed to all customers.

3. Data description

3.1 Water system data

We obtained characteristics of public water systems as well as violations from the EPA via a Freedom of Information Act (FOIA) request. The water system data covers the universe of U.S. public water systems and provides information such as population served, whether the system is operational, system deactivation date if applicable, system location, whether the system primarily uses surface water or groundwater, and the ownership structure (public or private). The violation data contains every violation recorded by the EPA since the agency began keeping records, using the Safe Drinking Water Information System (SDWIS).

Health-based water quality violations occur either because the concentration of a regulated contaminant exceeds a specified threshold (a maximum contaminant level, or MCL) or because a water system fails to follow an established technique for avoiding a contaminant violation (a treatment technique, or TT violation). The EPA also collects information on two additional types of violations; monitoring and reporting (MR) violations and “Other” violations. MR violations make up the vast majority of all violations and essentially amount to procedural violations. For example, if a utility fails to take a water quality reading at a specified time or does not monitor the quality of its water source, it incurs an MR violation. “Other” violations also involve procedural violations. Water systems that do not submit an annual water quality report, or fail to notify their customers of any acute health risks discovered in their water would incur violations classified as “Other.”

Audits conducted by EPA and external evaluators have established long-standing reporting problems in the SDWIS data (USGAO 2011). Recent work by the Government Accountability Office (GAO) reveals that the main problem seems to be one of under-reporting, primarily due to the failure of states to report some observed violations to the EPA at the national scale. This problem is more severe for MR violations than for health-based violations – GAO estimates that states correctly report 76 percent of health-based violations to the EPA, but only correctly report about 16 percent of MR violations (USGAO 2011).[8] Accurate reporting of microbial violations appears to exceed that for health-based violations overall (USGAO 2011, pp. 14-15).

Given that the vast majority of monitoring and reporting violations appear to go unreported (while the reporting rate for health-based violations is high), and the fact that health-based violations may have more important consequences, we focus on health-based water quality violations (and the subset of better-reported microbial violations) from community water systems in the 50 U.S. States.[9] Note that if violations are under-reported as audits suggest they may be, this will introduce measurement error in our dependent variable. If all systems under-report violations but the under-reporting is not systematically correlated with characteristics like system size or the propensity to violate, this measurement error will increase the standard errors in our models, making it more difficult to detect an effect of information disclosure. If reporting is systematically correlated with system characteristics, this will introduce bias in our estimates.[10]

Community water systems comprise the subset of public water systems that serve a non-transient community for more than six months out of the year.[11] As of July 2014 (the date of our FOIA request), community water systems made up about 34 percent of all active water systems, but served nearly 300 million persons. In contrast, our data indicate that non-community water systems served about 19 million persons.[12]

We match water system information with violations data, generating a panel of water quality violations in each system-year from 1990 to 2001 for systems that remained active throughout

[8] This difference may be due to the fact that EPA has established a goal that 90 percent of health-based drinking water violations be "completely and accurately reported to SDWIS/Fed," while it has no such goal for monitoring violations (USGAO 2011, p. 10).

[9] We exclude water systems managed by Native American tribes and U.S. territories.

[10] Prior work has evaluated the likelihood of strategic avoidance behavior in self-reported data among wastewater treatment plants and found little evidence that this is a problem (Shimshack and Ward 2005). In contrast, prior work has detected regulatory avoidance in reporting among public drinking water systems (Bennear et al. 2009). Given that we observe only violations in the SDWIS data, and not contaminant concentrations during and outside of periods in which violations occur, the data do not allow us to test for strategic reporting, except for the potential tradeoffs between microbial and DBP violations discussed in Section 4.3 and reported in Section 5.4.

[11] While the EPA defines public water systems as those systems serving at least 25 customers, our data indicate that many community water systems serve fewer than this threshold. Our main results change little if we exclude systems serving fewer than 25 customers.

[12] The EPA defines two types of non-community water systems: transient non-community water systems (such as gas stations and campgrounds), and non-transient, non-community water systems (such as schools, factories and hospitals). Both are excluded from our sample.

this time frame. Beyond 2001, changes in water quality standards for various contaminants complicate identifying an effect of the disclosure policy.[13] By keeping water systems that have remained active between 1990 and 2001, we also ensure that we have a balanced panel.[14]

The full panel includes 46,900 water systems observed over 12 years, for a total of 562,800 observations. Table 1 provides summary statistics for water systems in the full panel. The average water system serves 6,104 customers. Most water systems, however, are small. About 55 percent of systems (25,847 systems) serve 500 or fewer individuals. A further 36 percent (16,906 systems) of systems serve between 501 and 9,999 customers, 8 percent (3,732 systems) serve between 10,000 and 99,999 customers, and the final 1 percent (415 systems) serve at least 100,000 customers. Southern states (EPA region 4) account for the highest share of water systems. The average water system relies primarily on groundwater, does not purchase its water from another system, and is equally likely to be publicly or privately owned.

For each water system-year in our panel, we record the number of health-based violations incurred (Table 2). Health-based violations are infrequent, occurring in 6.9 percent of system-years. MR violations are much more common. Though the smaller systems generate the most violations, smaller systems are not appreciably more likely to incur a violation than larger systems. In fact, the systems most likely to incur a violation are those serving between 10,000 and 99,999 customers.

MCL violations account for 74 percent (47,452 violations) of health-based violations, with TT violations making up the remaining 26 percent (16,470 violations). Breaking down violations another way, almost 84 percent of health-based violations represent violations of microbial standards (53,398 violations). Violations associated with DBPs are rare in our sample.

Depending upon the specific nature of the violation and response of the water utility, violations can remain active for several months or even a year or more before the water system comes back into compliance. If a water system first reports a health-based violation to the EPA in March

[13] Our main concerns are the arsenic standard, which was updated in 2001 with an effective date of February 2002, and the Long Term I Enhanced Surface Water Treatment Rule along with new standards for disinfection byproducts and inorganic chemicals, which applied differentially by system size beginning in January 2002. These rules created compliance challenges, especially for small systems, starting in 2002 (Allaire et al. 2018).
[14] Some water systems report deactivation dates in the early 1900s. Since the EPA did not exist before 1970, we reason these dates could be a product of a Y2K bug, and drop 21 systems with deactivation dates prior to 1950.

1999 and comes back into compliance in January 2000, we count this health-based violation as a single violation recorded in 1999.[15]

3.2 Census and voting data

To examine the impact of service population demographics on the effects of disclosure, we merge our sample with demographic data from the U.S. Census (Adams 1990). We use the 1990 Census data, which predates our sample, in order to avoid any endogeneity concerns (for example, if low water quality prompts households to relocate). These data provide county-level information on mean family income and the percentage of the population that is White, Black, high school educated, educated with some college, and college educated. We assign Census variables to the county a water system operates in, making use of EPA's 2013 GPRA pivot tables, which report the county in which a system is located.[16] The average water system serves a population with a mean income of $38,142 (in 1990 dollars), and one that is about 87 percent White and one-third high-school educated (Table 1). On average, 17 percent of individuals have a college education.[17] As Table 1 shows, we are unable to match 5,135 systems (about 11 percent of our sample) with county names.

Finally, we obtain county-level data on voting in U.S. presidential elections during the study period from the Congressional Quarterly Press (CQ) voting and elections collection to examine whether water systems may respond differentially to information disclosure based on the political affiliation of their service population.[18] Because presidential elections occur once every

[15] How disclosure affects the duration of violations is an important question for future work.

[16] For large water systems that serve multiple counties, the GPRA table picks one county at random and assigns that county to the water system (pers. comm. K. Roland). While this may introduce some mismatches between water systems and household demographics, we believe the error introduced to be small, since only a few water systems are sufficiently large to span multiple counties.

[17] If we compare these average demographic characteristics for the community water systems in our data to the 1990 Census data for the whole United States, the population served by community water systems (those serving a non-transient population for at least 6 months per year) is slightly wealthier, slightly less Black and less White, and has somewhat higher maximum education than does the U.S. population as a whole. This population is also slightly more likely to reside in a county in which the Democratic candidate received a majority of votes in the most recent presidential election. These differences can be attributed to the fact that when we summarize demographics at the county level in our paper, a county enters our dataset once for each system located in that county. So where many smaller water systems reside within the same county, our average demographic variables over-weight that county, relative to the U.S. average which counts each county only once. In addition, areas not served by CWS (for example, rural areas where households have private groundwater wells) are included in the U.S. Census data, but not in our sample.

[18] https://library.cqpress.com/elections/download-data.php, last accessed April 29, 2022.

four years, we assign the 1992 election results to panel years 1990-1995, the 1996 election results to panel years 1996-1999, and the 2000 election results to panel years 2000-2001. Using these data, we calculate the share of votes for Democratic presidential candidates and then create an indicator variable set equal to 1 if the Democratic candidate garnered more than 50 percent of the sum of Democratic and Republican votes, else zero. As Table 1 indicates, we are unable to match 5,538 systems (about 12% of our sample) with county names in the political data, with many of the missing counties in Alaska, a state not included in the CQ voting data.

3.3 Pre-estimation matching

Ferraro and Miranda (2017) show that applying matching strategies prior to estimating panel fixed-effects models achieves results that more closely align with results obtained from randomized experiments. We follow this approach by creating three matched panels, one for each disclosure threshold. For the publishing requirement affecting systems serving over 500 customers, we draw potential matches from water systems serving at most 500 customers. For the mailing requirement affecting systems serving over 10,000 customers, we draw potential matches from water systems serving between 500 and 10,000 customers. For the web requirement affecting water systems serving over 100,000 customers, we draw potential matches from water systems serving between 10,000 and 100,000 customers. We match on system characteristics and pre-treatment health-based water quality violations, applying a 1-to-1 nearest neighbor matching algorithm (with replacement) using the Mahalanobis distance metric with exact matching on the EPA region, purchaser indicator, publicly-owned indicator, and surface water indicator. We further match on average pre-treatment health-based water quality violations. For each disclosure threshold, we generate a match for each treated water system. We randomly pick a control water system for systems having more than one match.[19] Table 3 illustrates average system characteristics of the matching variables for each disclosure threshold across treated and untreated water systems. Overall, Table 3 demonstrates that the matching successfully balances the average characteristics of our treated and untreated water systems. The mean pre-treatment violations for the publishing threshold is one exception: mean pre-treatment violations for water systems not required to publish equals 0.117, while mean pre-treatment

[19] Due to missing values of the matching variables, we are unable to find matches for 20 water systems treated with the publishing requirement and therefore drop them from our sample.

violations for water systems required to publish equals 0.127. Our differences-in-differences design, however, will account for pre-existing differences provided such differences remain constant over time. In our results below, we run placebo tests exploring this "parallel trends" assumption and find evidence in support of stable pre-trends.

4. Empirical models

4.1 Basic model

Our main model, Eq. (1), follows Bennear and Olmstead (2008). We regress v_{it}, the number of health-based violations by water system i in year t, on an interaction between a treatment indicator, T_i, and an indicator $Post_t$ that describes observations in or after 1998, a flexible function of system size, $\sum_{j=0}^{k} \delta_j (Post_t \times size_i)^j$, where *size* describes the service population in 100,000s, system fixed-effects (u_i,), and state-by-year fixed-effects (γ_{st}). The treatment indicator T_i equals one if the system serves a population exceeding one of the three disclosure thresholds; 501, 10,000, or 100,000 customers.[20]

$$v_{it} = \theta(T_i \times Post_t) + \sum_{j=0}^{J} \delta_j (Post_t \times size_i)^j + \gamma_{st} + u_i + \varepsilon_{it} \tag{1}$$

In Eq. (1) θ is the differences-in-differences (DID) estimate of the impact of requiring the respective disclosure method (publishing, mailing, or online posting) on health-based drinking water quality violations, net of any effect of system size away from the disclosure size threshold (see Online Appendix A for a mathematical derivation). We include the flexible function of system size to ensure that our treatment effect *(θ)* is not picking up size-related differences in the propensity to violate that are unrelated to the disclosure thresholds. For example, larger systems may have better administrative resources enabling compliance with water quality standards, or economies of scale in drinking water treatment may reduce compliance costs and make

[20] In our main specifications, we estimate a quadratic function of system size. As a robustness check, we vary J up to 6. Our overall conclusions remain unchanged.

compliance more likely (Bennear and Olmstead 2008, Raucher et al. 2011).[21] This makes our approach similar to a regression discontinuity (RD) design, though it is less dependent on observations very close to the size thresholds (Bennear and Olmstead 2008). In Online Appendix B, we estimate an RD model as an alternative to our DID-with-matching approach, generating qualitatively similar results.

Because systems below the information disclosure thresholds may voluntarily undertake more comprehensive disclosure than required, θ represents a lower bound on the effect of the information disclosure. For example, some systems serving fewer than 10,000 customers may mail their water quality reports. Any such costly, voluntary behavior by systems serving less than 10,000 customers would attenuate our estimate of θ, understating the true impact of mailing. The same argument also applies to the publishing and online posting requirements.

Given our approach, what are the main threats to identification? A key rationale for our use of the DID model that controls for system size away from the disclosure thresholds is that confounders must induce spurious changes in violations discontinuously at the disclosure thresholds *in 1998* (the first year for which CCRs report violations to consumers), or result in differences in violation trends among treated and control water systems that are not addressed by the independent variables in our model. The latter problem we test for directly in Section 5.5 with a series of placebo tests. The former could be a concern if there are other discontinuities at our system size thresholds in 1998. We find little evidence in the language of the 1996 Amendments or in the literature that supports this possibility, and we address this question carefully in Online Appendix C.[22]

4.2 Persistence of information disclosure

[21] For example, Raucher et al. (2011) show that the cost per household of complying with the updated arsenic MCL range from $407/household for the smallest systems to $1/household for the largest.

[22] The only component of the 1996 Amendments that has the potential to confound our estimates is the establishment of a small technical assistance fund available for systems serving between 25 and 10,000 individuals, which provided $15 million per year from 1997-2003 for all such systems, including many outside of our sample – tribal systems and those in the U.S. territories (Ramseur and Tiemann 2019). About 43,000 of our sample systems would have been eligible to apply for these funds; if divided equally across only our eligible sample systems (a conservative assumption), this would amount to less than $350 per system-year. If these funds went disproportionately to "frequent violators," our system fixed-effects may partially control for any effect of this technical assistance. Because we cannot control directly for receipt of these funds in our models, however, they are a potential, if unlikely, threat to identification. See Online Appendix C for additional detail on this issue.

At the outset, it is unclear whether we should expect the impact of the information disclosure policy to erode, persist, or strengthen over time. On the one hand, a service population may become accustomed to poor water quality reports and instead of voice concern, engage in avoidance behavior, such as bottled water purchases (Graff Zivin et al. 2011). The poorly-performing water system may then learn over time that poor reports have few consequences, reducing or eliminating any effect on violations. On the other hand, water quality reports may increase water system operators' attentiveness to violations, leading to increased expertise associated with reducing violations. Or they may make long-term investments (in treatment equipment, for example) that permanently reduce the likelihood of violations. This kind of persistence has been seen in the response of consumers to information treatments about energy and water conservation (Allcott and Rogers 2014, Ferraro et al. 2011). Either of these responses could lead to stable or even increasing impacts of the disclosure policy over time.[23]

To assess the persistence of impacts for each disclosure threshold, we estimate Eq. (2), interacting our treatment with annual dummy variables from 1998 through 2001. The estimates of β_t describe the impact of the information disclosure policy in each post-disclosure year.

$$v_{it} = \sum_{t=1998}^{2001} \beta_t (year_t \times T_i) + \delta f(size) + \gamma_{st} + u_i + \varepsilon_{it} \tag{2}$$

4.3 Heterogeneity in water system response

Responding to a growing environmental justice literature and prior work suggesting that responses to information disclosure may vary with income, education and other characteristics, we interact $(T_i \times Post_t)$ in Eq. (1) with a set of county-level demographic variables (D_c) from the U.S. Census, described earlier: mean household income, percent White, percent Black, percent with a high-school education, percent with some college education, and percent with a college degree. The demographic variables reflect the county in which each water system is located, given that we are unable to identify water system boundaries for the 46,900 systems in

[23] We distinguish between learning how to address violations and learning about the existence of violations. Water systems were reporting violations to the EPA for many years prior to 1998, so it seems unlikely that the information disclosure would cause utilities to learn about the existence of violations (Bennear and Olmstead 2008).

our data.[24] Similarly, we interact the Democratic share of presidential votes with $(T_i \times Post_t)$ to see if water systems may respond differentially to disclosure based on the environmental policy preferences of their service population. We estimate Eq. (3) separately for each demographic variable of interest.

$$v_{ict} = \mu(T_i \times Post_t \times D_c) + \theta(Post_t \times T_i) + \omega(Post_t \times D_c) + \sum_{j=0}^{J} \delta_j (Post_t \times size_i)^j + \gamma_{st} + u_i + \varepsilon_{it} \quad (3)$$

When testing for heterogeneous responses by political preference of water systems' customers, we modify Eq. (3) slightly. Unlike the county demographic variables, the Democratic share of presidential votes changes over time (with each election). Thus, the political heterogeneity regression equations also include both the Democratic share on its own, and an interaction between T_i and the Democratic share, in addition to the other variables on the right-hand side of Eq. (3).

We also test for heterogeneity in water system responses to information disclosure across violations for different drinking water contaminants. To explore the trade-off between microbial violations and DBP violations described earlier, we employ the seemingly unrelated regression framework shown in Eq. (4). All coefficients are analogous to those in Eq. (1), however each equation now refers to either microbial health-based violations, denoted by superscript *m*, or DBP violations,[25] denoted by superscript *dbp*. To estimate this system of equations, we follow the method suggested by Blackwell (2005).[26]

---

[24] We make the water system-county match using EPA's 2013 GPRA pivot tables. For large water systems that serve multiple counties, the GPRA table picks one county at random and assigns that county to the water system (pers. comm. K. Roland). While this possibly introduces mismatches between water systems and household demographics, we believe the error introduced to be small since only a few water systems are sufficiently large to span multiple counties. We are unable to match 5,135 systems (about 11 percent of our sample) with county names, and these systems are dropped from the analysis in this extension. We are unable to match 5,538 systems with county names in the political data, with many of the missing counties in Alaska, a state not included in the CQ voting data. These unmatched systems are dropped from the analysis in this extension.

[25] Trihalomethanes are the most relevant disinfection byproduct in our sample. Other DBPs such as bromate, chlorite, and haloacetic acids were not regulated until after 2001.

[26] Blackwell (2005) proposes defining the system of two equations as $y_1 = b_1 * \mathbf{X}_1 + e_1$, and $y_2 = b_2 * \mathbf{X}_2 + e_2$ (where y_1, y_2, b_1, b_2, e_1, and e_2 represent vectors, and $\mathbf{X}_1$ and $\mathbf{X}_2$ represent matrices of covariates). Blackwell suggests stacking the two equations in a diagonal matrix as follows $[y_1; y_2] = [\mathbf{X}_1\ 0; 0\ \mathbf{X}_2]\ [b_1; b_2] + [e_1; e_2]$, and running a regression on the newly formed single equation. Tests can then be performed on coefficients of interest. This is the procedure we implement.

$$v_{it}^{m} = \theta^{m}(T_i \times Post_t) + \delta^{m} f(size) + \gamma_{st}^{m} + u_i^{m} + \varepsilon_{it}^{m}$$
$$v_{it}^{dbp} = \theta^{dbp}(T_i \times Post_t) + \delta^{dbp} f(size) + \gamma_{st}^{dbp} + u_i^{dbp} + \varepsilon_{it}^{dbp} \tag{4}$$

Evidence of water systems reducing microbial contaminants at the expense of DBPs depends upon the signs and relative magnitudes of θ^{m} and θ^{dbp}. If water systems reduce both types of violations in response to the information disclosure requirement equally, we expect $\theta^{m} \approx \theta^{dbp} < 0$ (relative to their respective baseline levels of violations). If water systems reduce microbial violations at the expense of DBPs, and do so differentially in response to disclosure, we would expect $\theta^{m} < 0 < \theta^{dbp}$. If both coefficients indicate reductions in violations, but the magnitude of the effect on microbial violations is larger than that on DBP violations (for example, $\theta^{m} < \theta^{dbp} < 0$), then there may be simultaneously real reductions in water quality violations along with some trade-off of DBPs for fewer microbial contaminants.

5. Results

Figure 1 illustrates health-based violations in each year for the three matched panels described in Section 3, normalizing violations by the number of systems (in thousands) as well as the number of MCL regulations systems must comply with in each year of our panel. These normalizations account for the fact that many systems serve few customers, and that regulatory stringency has changed over time. Becoming effective at the end of 1990, both the Surface Water Treatment Rule and the Total Coliform Rule drive the early increase in normalized violations. Beginning in the early to mid-1990s, normalized violations fall for all categories of water systems, with the largest declines experienced by those water systems serving the most customers, and the smallest declines experienced by the systems serving the fewest customers. While the visual effects are subtle, Figure 1 provides some evidence for declining water violations beginning in 1998 for the publishing and web thresholds (relative to their control groups).

Table 4 reports numerically what Figure 1 illustrates graphically, showing the raw differences in means on either side of each information disclosure threshold using our three matched panels. At the publishing threshold, average violations for those systems serving over 500 customers fall by 0.023 violations after 1998. In contrast, average violations increase slightly for matched systems serving 500 or fewer customers after 1998, leading to a raw differences-in-differences estimate of -0.039 violations for systems above the publishing threshold. Raw differences in means also imply reductions in violations at the mailing and online posting thresholds. The empirical analysis that follows tests whether these patterns in the raw data are plausibly causal.

5.1 Main difference-in-difference results

Table 5 presents our main estimates of the impact of the publishing, mailing, and online posting requirements. We run all these and subsequent models with frequency weights to account for the fact that some control water systems were selected more than once in our matching procedure.

Columns 1 and 2 examine the impacts of the baseline requirement to publish a CCR on a water system's health-based violations. Column 1 includes only the linear system size variable, and column 2 adds a quadratic term.[27] The CCR publishing requirement reduces health-based violations by about 0.04 violations per system-year, corresponding to a 33 percent decrease in violations. [28] Notice that this basic result is very similar to the raw results in Table 4. Column 3 of Table 5 restricts the model to include only violations of microbial standards, with a similar result.[29]

Next, we test for an impact of the requirement to mail CCRs on all health-based (col. 4) and microbial (col. 5) violations, finding no significant impact. In columns 6 and 7, we do the same for the requirement to post the CCR on the web, and find a somewhat larger, statistically significant decrease in both health-based and microbial violations.

In columns 8 and 9, we report results from a dose-response model, incorporating all disclosure thresholds in one model to test for an additional incremental effect of each type of required reporting, relative to systems with no reporting requirement. These models use one general pre-

[27] While not shown, we include up to a sixth-order form of our flexible function of system size and find results robust to the higher order specifications for each disclosure threshold. These results are available from the authors on request.

[28] Average violation count equals 0.119 violations for systems subject to the publishing requirement (see Table 2), and our preferred estimate derives from the model incorporating the quadratic function of system size (column 2 of Table 5), so 0.039 / 0.119 = 0.328.

[29] Columns (3), (5) and (7) report estimates from just the microbial violations portion of Eq. (4). That is, the estimated model in each of these columns is in the same basic form as Eq. (1), but we use the SUR framework, incorporating both the microbial equation and the DBP equation in Eq. (4), to estimate the coefficients and standard errors. We do this to avoid reporting results from the same basic model in multiple locations in the paper. We discuss the full SUR model as an extension of our main model in Section 5.4, with the full results (including those for DBPs) reported in Online Appendix D.

estimation match for all treated systems (those serving >500), with matches drawn from the untreated group (those serving <500). In other words, we use the matched sample that we generated to explore the effect of the publishing threshold (see Table 3b). In these models, the baseline requirement to publish a CCR is similar to that in columns 1-3, the requirement to mail a copy of the CCR has a smaller additional violation-reducing effect, and the web posting requirement has no effect. Again, results are similar for all health-based violations in column 8, and for only microbial violations in column 9.

The two approaches in columns 1-7 and columns 8-9 differ in important ways. Thus, it is not surprising that the results differ somewhat. In columns 1-7, recall that we match each treatment group (the publishing group in columns 1-3, the mailing group in columns 4-5, and the online group in columns 6-7) with a set of control systems specific to that disclosure threshold, using the three different matched panels described in Section 3. In the mailing and online models, all control systems serve more than 500 individuals and are thus subject to some form of information disclosure. Like the publishing model in columns 1-3, in columns 8 and 9, none of the control systems are required to produce a CCR. On the one hand, the pre-estimation matching in columns 1-7 (especially the mailing and online models) likely provides a better set of controls for treated systems. On the other hand, the inclusion of all treatment groups in columns 8 and 9 allows us to test for the marginal effects of increasingly stringent disclosure requirements. Given these different approaches, we are not surprised that the online disclosure requirement appears to reduce violations for the very largest systems in columns 6 and 7, while mailing alone does not (in columns 4 and 5), while the reverse appears to be true when all three mechanisms are considered together in columns 8 and 9. However, given the consistent magnitude and significance of the basic CCR publishing requirement in row 1 of Table 5, many of our extensions and robustness checks going forward focus on this primary disclosure threshold.

5.2 Extension: persistence of information disclosure impacts

Figure 2 illustrates the persistence of the information disclosure policy, plotting results from estimating Eq. (2) for all health violations and microbial violations, and showing point estimates and confidence intervals for the impact of each disclosure requirement in 1998, 1999, 2000, and 2001. Across each year for both violation types, we observe relatively stable effects due to the publishing requirement and stable, statistically insignificant impacts over time for the mailing and online posting requirements.[30]

Figure 2 shows that the primary impact of the disclosure policy takes place directly after the policy, with the effect remaining relatively unchanged through the end of our sample. Figure D1 in Online Appendix D graphs estimates from a dose-response persistence model following our approach in columns 8-9 of Table 5, in which all three disclosure thresholds appear in the same equation, interacted with each year post-policy. Results are qualitatively similar to those in Figure 2. Taken together, the persistence results suggest that the steps taken by water systems to reduce violations when required to disclose compliance information annually to their customers may have been permanent (e.g., capital investments, or long-run changes in treatment and other management protocols).

5.3 Extension: response heterogeneity by system demographics and political preference

Table 6 reports results from our empirical tests of whether water systems respond differently to information disclosure requirements, depending on the demographics and political preference of the customers they serve. Given the robust, consistent results thus far for the primary CCR publishing requirement, we focus this analysis on heterogeneous responses to the fact of disclosure rather than the form.

Column 1 of Table 6 reports results from our main model (comparable to column 2 of Table 5), using the sub-sample of system-years we were able to match with county-level demographic data from the U.S. Census. Columns 2-7 report results from estimating Eq. (3), using the baseline publishing requirement as the treatment, and interacting the treatment (and, for completeness, the post-treatment dummy) with six different summary demographic variables: mean household income, percent White, percent Black, percent with a high-school education, percent with some college, and percent with a college degree.

---

[30] We also estimate a model that interacts our treatment indicator with an annual linear time trend describing the number of years since 1998 and come to similar conclusions. These results stand in contrast to Fung et al. (2007), who postulate that the information disclosure policy we study here will have less of an effect over time on account of few disclosure-associated benefits realized by water systems and the absence of citizen groups to give voice to concerns over water quality.

The average income of households in the county in which a water system is located does not appear to impact the effect of disclosure on violations (the triple interaction term, μ from Eq. (3), is not significantly different from zero). The standard error on our main treatment effect in column 2 is large, making this effect statistically insignificant. Similarly, the triple interactions with the two racial composition variables (*%White* and *%Black*) are not significantly different from zero.

While the interaction with *%HS educ* is similarly insignificant (column 5), the result in column 6 suggests that there may be some difference in water systems' propensity to reduce violations across systems with differing levels of education. In these models, the baseline effect of disclosure in the first row is somewhat larger than the result in column 1. The higher is the presence of college education among the served population, the larger (weakly) is the reduction in violations in the post-policy period. However, the triple interaction estimate – μ from Eq. (3) – in column 6 is not terribly informative. In column 6, a higher percent of the served population with some college weakly shrinks the violation-reducing impact of disclosure, though the effect is very small.

The results in column 8 include our interaction between the treatment effect and a variable indicating whether a water system's county had a Democratic majority in the most recent U.S. presidential election. Similar to the demographic results described above, we find no statistically significant effect of this political variable on the impact of disclosure on violations, suggesting that water systems do not respond differentially to disclosure based on the political leanings of their service areas. Note that, because the presidential election vote shares vary over time, column 8 includes the Democratic vote share variable on its own, as well as an interaction with T_{pub}, neither of which is significant.

Prior work has found that some information disclosure policies can induce heterogeneous responses by income and education (Powers et al. 2011, Shimshack et al. 2007). Reasons for our different results could stem from the fact that Powers et al. (2011) focus on a very different context -- they show that, in response to a green rating program, pulp and paper plants in India reduce pollution loadings more in wealthier communities than they do in poorer communities and find no effect of variation in communities' literacy rate. The regulated entities are private firms, and the effects of the ambient pollution disclosed by the policy are diffuse (rather than directly affecting households' drinking water). Shimshack et al. (2007) find that education and newspaper readership are important predictors of *consumer* responses to mercury warnings for canned fish, again, a different context from our own. Banzhaf et al. (2019) note that the assumption that the effect of exposure to an environmental hazard coincides with people living in the same geographic unit as the hazard ("unit-hazard coincidence") is problematic, partly because it can introduce measurement error that attenuates estimates of environmental justice correlations. In our case, additional measurement error is introduced by the fact that we are unable to observe actual water system boundaries for the more than 46,000 systems in our sample, so our summary statistics on racial and socioeconomic composition of households in the county in which a system is located may not accurately represent the characteristics of households served by each system. However, other papers in the economics literature fail to find that the impacts of environmental policies are causally linked with demographic characteristics of the affected households (Fowlie et al. 2012, Mansur and Sheriff 2021).

5.4 Extension: Potential tradeoffs between microbial and DBP violations

The next extension of our main models explores potential tradeoffs between microbial violations and DBP violations that would occur if water systems respond to the information disclosure requirement by increasing disinfection so as to reduce the likelihood of more salient microbial violations (Adamowicz et al. 2011), while potentially increasing the likelihood of less salient DBP violations.

Columns (3), (5) and (7) of Table 5 report results from the microbial portion of the two-equation system (Eq. 4) that we estimate to conduct these tests. The SUR framework also estimates coefficients for the DBP portion of Eq. (4), but none of the DBP model coefficient estimates were significantly different from zero in any of the SUR models we estimated. Following the estimation of each SUR model, we test for statistical differences between the estimated DBP and microbial treatment effect coefficients (θ^{m} and θ^{dbp} in Eq. 4). The estimated effects do differ statistically for the baseline CCR publishing requirement (F=9.02 and p=0.0027) and for the online posting requirement (F=6.48 and p=0.0110), but not for the mailing requirement (F=1.85 and p=0.1741). For completeness, we report the full SUR results, including the DBP results, in Table D1 of Online Appendix D.

As noted earlier, DBP violations are very infrequent, so our null results may be attributable to the small amount of variation in the dependent variable. Nonetheless, our sample includes all community water systems in the United States operating during this time period (that is, there is no larger sample with additional DBP violations available), so these results are consistent with the conclusion that water systems reduced microbial violations in response to the CCR publishing requirement (column 3 of Table 5), without concurrently increasing DBP violations.

Water systems have several means by which to reduce microbial contaminants. Trihalomethanes arise from the chlorination process. Disinfecting water using ozonation or chlorine dioxide produces other DBPs (bromate and chlorite, respectively), which were unregulated during our sample period. Water systems may have switched disinfection methods to avoid violating the DBP standard (which primarily affected trihalomethanes) while maintaining the microbial standards; these additional DBPs were regulated in 2002, so even if this behavior (unobservable during our sample period) was a potential concern, it was subsequently addressed with new regulation.[31]

5.5 Robustness checks

The validity of our identification strategy relies on the counterfactual assumption that treated systems would have behaved similarly to non-treated systems in the absence of the disclosure policy. To examine the validity of this assumption, we conduct some placebo tests, imposing a proxy information disclosure policy in 1994, 1995, 1996, and 1997 (when there was none in place), and run our model on the pre-period sample. Tables D2-D5 in Online Appendix D show statistically insignificant coefficient estimates on the interaction between the treatment indicator and the placebo policy indicator for each disclosure threshold in 1994, 1995, 1996, and 1997. These estimates provide supportive evidence for parallel trends in the pre-period data. The top panel of Figure 1 also provides some evidence of parallel trends among non-disclosing and disclosing water systems in the pre-policy period. (Recall that the SDWA Amendments were

[31] The literature does present evidence of strategic regulatory avoidance behavior on the part of water systems with respect to microbial violations (Bennear et al. 2009). However, we cannot test comprehensively for this behavior with respect to CCRs, because the “substitute” DBPs – those that might have been elevated if systems switched disinfection methods to avoid microbial violations – were not monitored until they were regulated in 2002.

passed in 1996, and that the first CCRs were issued in July 1999, covering violations from the 1998 calendar year.)

We observe water systems' service population in 2014, because there are no historical, time-varying system characteristics data available from EPA. Since we assign water systems to disclosure treatment using 2014 system size, if water systems experienced substantial population change between 1998 and 2014, we may inadvertently assign some water systems to the wrong treatment group. To test for sensitivity to possible treatment misassignment, we run models that exclude systems with 2014 service populations within 10 percent, 20 percent, and 30 percent of the publishing threshold service population cutoff, reasoning that any water systems assigned to the wrong treatment status would be concentrated within these bands around the treatment threshold. Results in Table D6 in Online Appendix D show that the estimated coefficient on $T_{pub} \times Post$ is robust to these tests. Any misassignment of treatment does not appear to substantially impact our main results. Note that the number of observations in Table D6 changes little from column to column, suggesting that the number of systems "close" to the CCR publishing threshold is small, an additional reason to prefer the DID approach to RD.

Our disclosure treatment indicators only determine whether a water system serves above or below a certain population threshold and as a consequence, the matched panels developed to assess the effect of the publishing and mailing requirement do not distinguish whether a treated or control system may be subject to multiple disclosure requirements. Some water systems assigned to the publication treatment (i.e. those systems serving greater than 500 customers) also face the mailing or additionally the online posting requirements, and some water systems assigned to the mailing requirement additionally face the online posting requirement. To address any concern that including the full sample of water systems biases our results, we run two additional models designed to isolate the effect of the publishing and mailing disclosure thresholds.[32]

We present results in Online Appendix D, Table D7. In columns 1 and 2, we isolate the impact of the publishing requirement by limiting the sample to those water systems that serve fewer than

[32] Note that this is not a concern for the online posting requirement since any water system facing this requirement does not face any further requirement.

10,000 customers. For this model, a water system is considered treated if it served more than 500 customers and therefore must publish, and only publish, the CCR. In columns 3 and 4, we isolate the impact of the mailing requirement by restricting the sample to those water systems serving greater than 500 and less than 100,000 customers. For this model, the comparison is between "treated" water systems that serve between 10,000 and 100,000 customers and thus must mail, and only mail, the CCR and those systems that serve more than 500 but less than 10,000 customers and thus only must publish the CCR (and do not need to mail it). It is important to note that the definition of what is "treated" varies throughout this table as we seek to isolate the marginal effect of these two levels of disclosure.

In column 1, the result is very similar to our main model in column 1 of Table 5. When we limit the sample to those systems serving fewer than 10,000 customers, and also include the quadratic function of system size, however, we cannot identify an effect of disclosure on violations in this group, alone (column 2). The flexible function of system size is arguably less important in this system-size-limited sample, however, we note that our main result is not robust to dropping the 4,147 water systems serving more than 10,000 customers. That may suggest that larger systems are important drivers of the responsiveness we observe to the publishing requirement. In columns 3 and 4, we find no significant effect of the mailing requirement on violations, consistent with our main results in Table 5.

6. Conclusions

The 1996 SDWA Amendments require community water systems to provide annual water quality reports to their customers, using methods of disclosure that vary with system size. We show that water systems required to disclose reports in some fashion (publish in a local venue, mail, or mail and post the report online) reduce water quality violations by about 33 percent – a national estimate consistent with the findings for Massachusetts, alone, in Bennear and Olmstead (2008). Our results for more stringent reporting requirements (mailing reports directly to customers, and providing the reports online) are mixed. When we pool all three requirements into a single model, results suggest there may be some incremental violation reductions from the requirement to mail reports (but no additional incremental effect from online posting), while models that test each method separately using a matching approach specific to each reporting group suggest the opposite. Overall, our results may suggest that the fact of information disclosure, rather than the format, drives water systems' response.

The effects we estimate appear to persist for at least three years post-policy. Ours is the first paper to demonstrate a sustained effect of environmental information disclosure required of regulated entities. As such policies become more common, assessments of their long-run effects are critical to understanding their proper role in regulating environmental hazards.

We find no consistent evidence that water systems' response to the information disclosure requirement varies with the race, education, income, or political affiliation of the population in the counties where systems are located. Prior work has raised concerns that low-income or less-educated households may benefit less from environmental information disclosure than high-income or highly-educated households. Our results do not suggest that this is the case for disclosure regarding local water system compliance with federal drinking water regulations. Note that minority and low-income communities do bear a disproportionate burden from drinking water violations, however (Allaire et al. 2018). Even though information disclosure does not appear to have exacerbated this differential burden, it has not corrected it, either.

Returning to the question of mechanisms raised in our introduction, our results are most consistent with water system operators learning and making process improvements as a result of the requirement to disclose violations to their consumers each year, rather than political or market pressure. The market mechanism is unlikely given the small share of piped water consumed as drinking water and the insensitivity of our estimated effects to service area household income. Political pressure, while possible, is inconsistent with a lack of differential response based on service population demographics and political preferences.

Finally, if process improvements are driving our results, it is useful to consider whether water systems made any strategic tradeoffs in order to reduce violations. While we cannot rule out all strategic regulatory avoidance behavior, despite the fact that there is a known trade-off between reducing microbial contaminants and increased levels of carcinogenic disinfection byproducts, we have shown that water systems were able to reduce microbial contaminant violations without experiencing increases in DBP violations during our sample period. Among all possible DBPs, however, the EPA only regulated trihalomethanes throughout the timeframe of our analysis. Further work would be needed to determine whether water systems are incurring additional violations of less-salient DBP standards in order to reduce violations of more-salient microbial standards under today's set of MCLs.

Collectively, the empirical tests we perform show for the first time that the requirement that U.S. water systems disclose violations to end-users on an annual basis reduced health-based SDWA violations at the *national* scale by about one-third, that this impact did not appear to vary with water consumers' demographics, that the impacts of disclosure persisted for at least three years after the policy was implemented, and that it did not prompt one potential risk-risk tradeoff (of DBP for microbial violations). Debate over the proper structure of this policy continues. As part of a retrospective review in 2012, since 2013 the EPA has allowed systems to post CCRs online in lieu of mailing them (recall that the online posting requirement for the largest systems in the policy we evaluate was in addition to – not instead of – mailing CCRs). The bipartisan America's Water Infrastructure Act of 2018 requires that EPA promulgate a new rule mandating that CCRs be issued at least twice per year by systems serving more than 10,000 individuals, rather than once per year (though EPA has not yet promulgated a rule requiring this). A recent survey of U.S. water system managers indicates significant concern about the additional costs of this more frequent disclosure requirement (Evans and Carpenter 2019). Further work is needed to understand the mechanisms driving water system responses to information disclosure requirements under the SDWA, and to assess the benefits and costs of these rules.

References

Adamowicz, Wiktor, Diane Dupont, Alan Krupnick, and Jing Zhang. 2011. Valuation of cancer and microbial disease risk reductions in municipal drinking water. *Journal of Environmental Economics and Management* 61: 213-226.

Adams, Terry K. Census of Population and Housing, 1990 [United States]: Extract Data. Inter-university Consortium for Political and Social Research [distributor], 2006-01-18. https://doi.org/10.3886/ICPSR02889.v1

Allaire, Maura, Haowei Wu, and Upmanu Lall. 2018. National trends in drinking water quality violations. *Proceedings of the National Academy of Sciences* 115(9): 2078-2083.

Allcott, Hunt, and Todd Rogers. 2014. The short-run and long-run effects of behavioral interventions: experimental evidence from energy conservation. *American Economic Review* 104(10): 3003-37.

Allcott, Hunt, and Richard L. Sweeney. 2016. The role of sales agents in information disclosure: evidence from a field experiment. *Management Science* 63(1): 21-39.

Balasubramanya, Soumya H., Alexander Pfaff, Lori S. Bennear, Alessandro Tarozzi, Kazi Matin Ahmed, Amy Schoenfeld, and Alexander van Geen. 2014. Evolution of households' responses to the groundwater arsenic crisis in Bangladesh: information on environmental health risks can have increasing behavioral impact over time. *Environment and Development Economics* 19(5): 631-647.

Banzhaf, Spencer, Lala Ma, and Christopher Timmins. 2019. Environmental justice: the economics of race, place and pollution. *Journal of Economic Perspectives* 33(1): 185-208.

Bennear, Lori S., Katrina K. Jessoe, and Sheila M. Olmstead. 2009. Sampling out: regulatory avoidance and the total coliform rule. *Environmental Science & Technology* 43(14): 5176–5182.

Bennear, Lori S., and Sheila 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(2): 117–130.

Blackwell, J. Lloyd, III. 2005. Estimation and testing of fixed-effect panel-data systems. *The Stata Journal* 5(2): 202–207.

Delmas, Magali, Maria J. Montes-Sancho, and Jay P. Shimshack. 2010. Information disclosure policies: evidence from the electricity industry. *Economic Inquiry* 48(2): 483–498.

Dranove, David, and Ginger Zhe Jin. 2010. Quality disclosure and certification: theory and practice. *Journal of Economic Literature* 48(4): 935-963.

Earnhart, Dietrich. 2004. The effects of community characteristics on polluter compliance levels. *Land Economics* 80(3): 408-432.

Earnhart, Dietrich, and Paul J. Ferraro. 2021. The effect of peer comparisons on polluters: a randomized field experiment among wastewater dischargers. *Environmental and Resource Economics* 79: 627–652.

Evans, Jessica, and Adam T. Carpenter. 2019. Utility approaches to evaluating the effectiveness of consumer confidence reports. *Utilities Policy* 58: 136-144.

Ferraro, Paul J., and Juan José Miranda. 2017. Panel data designs and estimators as substitutes for randomized controlled trials in the evaluation of public programs. *Journal of the Association of Environmental and Resource Economists* 4(1): 281–317.

Ferraro, Paul J., Juan José Miranda, and Michael K. Price. 2011. The persistence of treatment effects with norm-based policy instruments: evidence from a randomized environmental policy experiment. *American Economic Review: Papers and Proceedings* 101(3): 318-322.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur. 2012. What do emissions markets deliver and to whom? Evidence from Southern California's NOx trading program. *American Economic Review* 102(2): 965-993.

Fung, Archon, Mary Graham, and David Weil. 2007. *Full Disclosure: The Perils and Promise of Transparency*. Cambridge, UK: Cambridge University Press.

Fung, Archon, and Dara O'Rourke. 2000. Reinventing environmental regulation from the grassroots up: explaining and expanding the success of the Toxics Release Inventory. *Environmental Management* 25: 115–127.

Graff Zivin, Joshua, Matthew Neidell, and Wolfram Schlenker. 2011. Water quality violations and avoidance behavior: evidence from bottled water consumption. *American Economic Review: Papers and Proceedings* 101(3): 448–453.

Karkkainen, Bradley C. 2001. Information as environmental regulation: TRI and performance benchmarking, precursor to a new paradigm? *Georgetown Law Journal* 89: 257.

Mansur, Erin T., and Glenn Sheriff. 2021. On the measurement of environmental inequality: ranking emissions distributions generated by different policy instruments. *Journal of the Association of Environmental and Resource Economists* 8(4): 721-758.

Mastromonaco, Ralph. 2015. Do environmental right-to-know laws affect markets? Capitalization of information in the toxic release inventory. *Journal of Environmental Economics and Management* 71: 54–70.

Powers, Nicholas, Allen Blackman, Thomas P. Lyon, and Urvashi Narain. 2011. Does disclosure reduce pollution? Evidence from India's Green Rating Project. *Environmental and Resource Economics* 50(1): 131–155.

Ramseur, Jonathan L., and Mary Tiemann. 2019. Water infrastructure financing: history of EPA appropriations. CRS Report 96-647, Washington, DC: Congressional Research Service. https://www.everycrsreport.com/files/20190410_96-647_981dcf4b10a712aad371e7b8bce38a4430f1608d.pdf

Raucher, Robert S., Scott J. Rubin, Douglas Crawford-Brown, and Megan M. Lawson. 2011. Benefit-cost analysis for drinking water standards: efficiency, equity, and affordability considerations in small communities. *Journal of Benefit-Cost Analysis* 2(1): 1-24.

Reinhart, Forest L. 2000. *Down to Earth: Applying Business Principles to Environmental Management*. Boston: Harvard Business School Press.

Shimshack, Jay P., and Michael B. Ward. 2005. Regulator reputation, enforcement, and environmental compliance. *Journal of Environmental Economics and Management* 50: 519-540.

Shimshack, Jay P., Michael B. Ward, and Timothy K. Beatty. 2007. Mercury advisories: information, education, and fish consumption. *Journal of Environmental Economics and Management* 53(2): 158–179.

Sunstein, Cass. 1999. Informational regulation and informational standing: Akins and beyond, *University of Pennsylvania Law Review* 147: 613–675.

Tiemann, Mary. 2017. Safe Drinking Water Act (SDWA): A Summary of the Act and Its Major Requirements. CRS Report RL31243, Congressional Research Service. https://fas.org/sgp/crs/misc/RL31243.pdf

U.S. General Accounting Office. 2011. Drinking Water: Unreliable state data limit EPA's ability to target enforcement priorities and communicate water systems' performance. GAP Report to Congressional Requesters, GAO-11-381. https://www.gao.gov/products/gao-11-381

Water Research Foundation. 2016. Residential End Uses of Water Version 2: Executive Report. https://www.waterrf.org/resource/residential-end-uses-water-version-2-executive-report.

Table 1: Summary statistics of water system characteristics

	Mean	SD	Min	Max	Obs.
Size (Service Population)					
All systems	6,104	59,471	0	8,271,000	46,900
Less than 501	160	132	0	500	25,847
501 to 9,999	2,630	2,288	501	9,999	16,906
10,000 to 99,999	28,285	19,684	10,000	99,750	3,732
Greater than 100,000	318,394	540,420	100,000	8,271,000	415
Greater than 500 (indicator)	0.45	0.50	0	1	46,900
At least 10,000 (indicator)	0.09	0.28	0	1	46,900
At least 100,000 (indicator)	0.01	0.09	0	1	46,900
EPA Region					
Region 1	0.05	0.22	0	1	46,900
Region 2	0.07	0.25	0	1	46,900
Region 3	0.09	0.29	0	1	46,900
Region 4	0.18	0.38	0	1	46,900
Region 5	0.15	0.35	0	1	46,900
Region 6	0.16	0.37	0	1	46,900
Region 7	0.08	0.27	0	1	46,900
Region 8	0.06	0.24	0	1	46,900
Region 9	0.09	0.29	0	1	46,900
Region 10	0.08	0.28	0	1	46,900
Other Water System Characteristics					
Water source (1 = surface)	0.23	0.42	0	1	46,869
Publicly owned systems	0.50	0.50	0	1	46,661
Systems purchasing water	0.18	0.38	0	1	46,869
Demographic & Election Result variables					
Service population income ($1990)	$38,142	$9,527	$15,363	$81,846	41,765
Percent white	0.87	0.14	0.05	1.00	41,765
Percent black	0.08	0.12	0.00	0.86	41,765
Percent high school education	0.33	0.06	0.14	0.53	41,765
Percent some college	0.24	0.06	0.07	0.45	41,765
Percent college	0.17	0.07	0.04	0.53	41,765
Percent Democrat wins	0.46	0.44	0.00	1.00	41,362

Notes: EPA region 1 (Connecticut, Massachusetts, Maine, New Hampshire, Rhode Island, Vermont), EPA region 2 (New Jersey, New York), EPA region 3 (District of Columbia, Delaware, Maryland, Pennsylvania, Virginia, West Virginia), EPA region 4 (Alabama, Florida, Georgia, Kentucky, Mississippi, North Carolina, South Carolina, Tennessee), EPA region 5 (Illinois, Indiana, Michigan, Minnesota, Ohio, Wisconsin), EPA region 6 (Arkansas, Louisiana, New Mexico, Oklahoma, Texas), EPA region 7 (Iowa, Kansas, Missouri, Nebraska), EPA region 8 (Colorado, Montana, North Dakota, South Dakota, Utah, Wyoming), EPA region 9 (Arizona, California, Hawaii, Nevada), EPA region 10 (Alaska, Idaho, Oregon, Washington). Water systems source water from either surface water or groundwater, and are either publicly owned by local, state, or federal government, or owned by a private entity (Native American entities own less than half a percent of these “private” water systems, and EPA classifies a further 5.6 percent of private water systems in our sample as owned by a mix of private and public entities).

Table 2: Summary of violations for system-years (unmatched panel)

	Totals	Mean	SD	Min	Max	$\Pr(v_{it} > 0)$	Obs.
Health violations by system service population							
All health violations	63,922	0.114	0.59	0	29	0.069	562,800
Systems serving 500 customers or less	33,933	0.109	0.55	0	21	0.067	310,164
Systems serving 501 to 9,999 customers	23,350	0.115	0.63	0	29	0.067	202,872
Systems serving 10,000 to 99,999 customers	6,244	0.139	0.63	0	24	0.087	44,784
Systems serving at least 100,000 customers	395	0.079	0.42	0	10	0.053	4,980
Systems serving over 500 customers	29,989	0.119	0.63	0	29	0.071	252,636
Systems serving at least 10,000 customers	6,639	0.133	0.61	0	24	0.084	49,764
Violations by violation type							
All violations	575,597	1.023	6.80	0	1,247	0.228	562,800
MCL violations	47,452	0.084	0.40	0	17	0.058	562,800
TT violations	16,470	0.029	0.42	0	29	0.011	562,800
Microbial violations	53,398	0.095	0.54	0	29	0.058	562,800
DBP violations	642	0.001	0.05	0	8	0.001	562,800
MR violations	485,282	0.862	6.71	0	1,247	0.167	562,800
Other violations	26,393	0.047	0.52	0	59	0.032	562,800

Notes: MCL violations refer to violations in which a regulated substance exceeds a maximum contaminant level specified by EPA. TT violations refer to treatment technique violations in which a water system fails to control a regulated substance in the manner specified by EPA. In our analysis, we consider MCL and TT violations to be health-based violations. We define two further subsets of health-based violations; microbial and DBP, or disinfectant by-product, violations. We define microbial violations as violations of regulations governing bacterial contaminants such as coliforms and giardia. We define DBP violations as violations of regulations governing disinfectant by-products. MR violations refer to violations of EPA's monitoring and reporting standards, such as a failure to report contaminant level testing results by a certain time. Finally, other violations refer to all other types of water quality violations. Failure to submit the annual water quality report represents one example of an 'Other' violation.

Table 3: Mean system characteristics for treated and untreated water systems pre- (a) and post-matching (b)

(a) Unmatched Panel		*Publishing*			*Mailing*			*Online*		
	Full	$T = 0$	$T = 1$	p-value	$T = 0$	$T = 1$	p-value	$T = 0$	$T = 1$	p-value
EPA Region	5.52	5.68	5.32	0.00	5.55	5.21	0.00	5.52	5.62	0.39
Water Source (1 = surface)	0.23	0.11	0.38	0.00	0.19	0.63	0.00	0.23	0.84	0.00
Publicly-owned systems	0.50	0.26	0.80	0.00	0.47	0.87	0.00	0.50	0.85	0.00
Systems purchasing water	0.18	0.12	0.26	0.00	0.17	0.31	0.00	0.18	0.24	0.00
Mean Pre-Treatment Violations	0.116	0.107	0.127	0.000	0.112	0.149	0.000	0.116	0.098	0.294
Observations	46,900	25,847	21,053		42,753	4,147		46,485	415	
(b) Matched Panel		*Publishing*			*Mailing*			*Online*		
	Full	$T = 0$	$T = 1$	p-value	$T = 0$	$T = 1$	p-value	$T = 0$	$T = 1$	p-value
EPA Region	5.52	5.31	5.31	1.00	5.21	5.21	1.00	5.62	5.62	1.00
Water Source (1 = surface)	0.23	0.38	0.38	1.00	0.63	0.63	1.00	0.84	0.84	1.00
Publicly-owned systems	0.50	0.80	0.80	1.00	0.87	0.87	1.00	0.85	0.85	1.00
Systems purchasing water	0.18	0.26	0.26	1.00	0.31	0.31	1.00	0.24	0.24	1.00
Mean Pre-Treatment Violations	0.116	0.117	0.127	0.007	0.141	0.149	0.315	0.092	0.098	0.717
Observations	46,900	21,033	21,033		4,147	4,147		415	415	

Table 4: Raw differences in means for each disclosure threshold

	Above Threshold			*Below Threshold*			
	Post-1998	Pre-1998	Δ_1	Post-1998	Pre-1998	Δ_2	$\Delta 1 - \Delta 2$
Publishing	0.103	0.127	-0.023	0.132	0.117	0.015	-0.039
Mailing	0.101	0.149	-0.048	0.102	0.141	-0.039	-0.009
Online	0.042	0.098	-0.056	0.092	0.092	0.000	-0.055

Table 5: Main results on the impact of the publishing, mailing, and online information disclosure requirements

	Health-based	Health-based	Microbial	Health-based	Microbial	Health-based	Microbial	Health-based	Microbial
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$T_{pub} \times Post$	-0.039	-0.039	-0.034					-0.036	-0.031
	(0.012)***	(0.012)***	(0.011)***					(0.012)***	(0.011)***
$T_{mail} \times Post$				-0.004	-0.012			-0.018	-0.020
				(0.013)	(0.011)			(0.007)**	(0.007)***
$T_{web} \times Post$						-0.059	-0.057	0.001	0.001
						(0.025)**	(0.021)***	(0.021)	(0.019)
$f(size)$	-0.006	-0.007	-0.006	-0.005	-0.002	-0.002	-8.2E-04	-0.003	-0.001
	(0.001)***	(0.003)***	(0.003)**	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
$f(size)^2$		3.5E-05	2.0E-05	-3.2E-06	-3.9E-05	-4.2E-05	-5.5E-05	-2.4E-05	-4.8E-05
		(0.000)	(3.8E-05)	(4.0E-05)	(3.8E-05)	(4.8E-05)	(4.4E-05)	(4.7E-05)	(4.4E-05)
Threshold	501	501	501	10k	10k	100k	100k	Dose-Resp.	Dose-Resp.
Fixed effects									
u_i	yes	yes	yes	yes	yes	yes	yes	yes	yes
γ_{st}	yes	yes	yes	yes	yes	yes	yes	yes	yes
adj. R^2	0.28	0.28	0.29	0.24	0.25	0.14	0.14	0.28	0.28
Systems	29,756	29,756	29,756	7,341	7,341	782	782	29,756	29,756
Observations	504,792	504,792	1,009,584	99,516	199,032	9,912	19,824	504,792	504,792

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors (reported in parentheses) are robust and clustered by water system. Water system fixed-effects are abbreviated u_i, and state-by-year fixed effects are γ_{st}. Threshold service populations listed are for the publishing (501), mailing (10k), and online disclosure requirements (100k). Columns (1), (2), (4), and (6) report estimates using Eq. (1). Columns (3), (5) and (7) report estimates from just the microbial violations portion of Eq. (4), with the full results (including those for DBPs) reported in Online Appendix B. Columns (8) and (9) we include all three reporting thresholds in the same equation.

Table 6: Tests for differences in disclosure responsiveness as a function of customer demographic and political characteristics

	No Dem	Income	%White	%Black	%HS educ	%Some coll	%Coll deg	Dvote>50%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
T_{pub} × *Post*	-0.035	-0.065	-0.070	-0.041	0.050	-0.121	-0.061	-0.039
	(0.010)***	(0.040)	(0.054)	(0.013)***	(0.057)	(0.052)**	(0.026)**	(0.012)***
Post × *Dem*		0.000	-0.001	-0.001	0.002	-0.003	-0.002	-0.019
		(0.000)	(0.001)	(0.001)	(0.002)	(0.002)*	(0.001)*	(0.018)
T_{pub} × *Post* × *Dem*		8.5E-07	4.1E-04	8.0E-04	-2.6E-03	3.7E-03	1.7E-03	0.014
		(9.3E-07)	(6.7E-04)	(5.5E-04)	(0.002)	(0.002)*	(0.001)	(0.019)
T_{pub} × *Dem*								-0.010
								(0.017)
Dem								0.002
								(0.017)
f(*size*)	-0.008	-0.008	-0.009	-0.008	-0.010	-0.009	-0.008	-0.007
	(0.003)***	(0.003)***	(0.003)***	(0.003)***	(0.003)***	(0.003)***	(0.003)***	(0.003)***
f(*size*)^2	4.4E-05	4.2E-05	4.7E-05	4.6E-05	6.2E-05	5.4E-05	4.5E-05	3.51E-05
	(3.9E-05)	(3.9E-05)	(4.0E-05)	(4.0E-05)	(4.1E-05)	(4.0E-05)	(3.9E-05)	(3.9E-05)
Characteristic	NA	Income	% White	% Black	% high school grad	% some college	% college degree	Democratic vote>50%
adj. R^2	0.27	0.27	0.27	0.27	0.27	0.27	0.27	0.25
Systems	27,679	27,679	27,679	27,679	27,679	27,679	27,679	27,476
Observations	458,928	458,928	458,928	458,928	458,928	458,928	458,928	455,208

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ Standard errors (reported in parentheses) are robust and clustered by water system. All models include water system and state-by-year fixed effects and are run for the publishing disclosure requirement threshold. Column (1) reports results from estimating Eq. (1), but on the sample of system-years that can be matched with the U.S. Census data. Columns (2) through (7) each include a triple-difference interaction between the treatment effect and one demographic/political characteristic of households in the county in which the water system is located. Column 8 represents a similar model, but with a variable indicating whether a water system's county had a Democratic majority in the most recent U.S. presidential election (note that because the presidential election vote shares vary over time, Column 8 additionally includes the Democratic vote share variable on its own, as well as an interaction with *Tpub*, neither of which is significant).

Figure 1: Health-based water quality violations by year, comparing treatment and control groups for each disclosure threshold

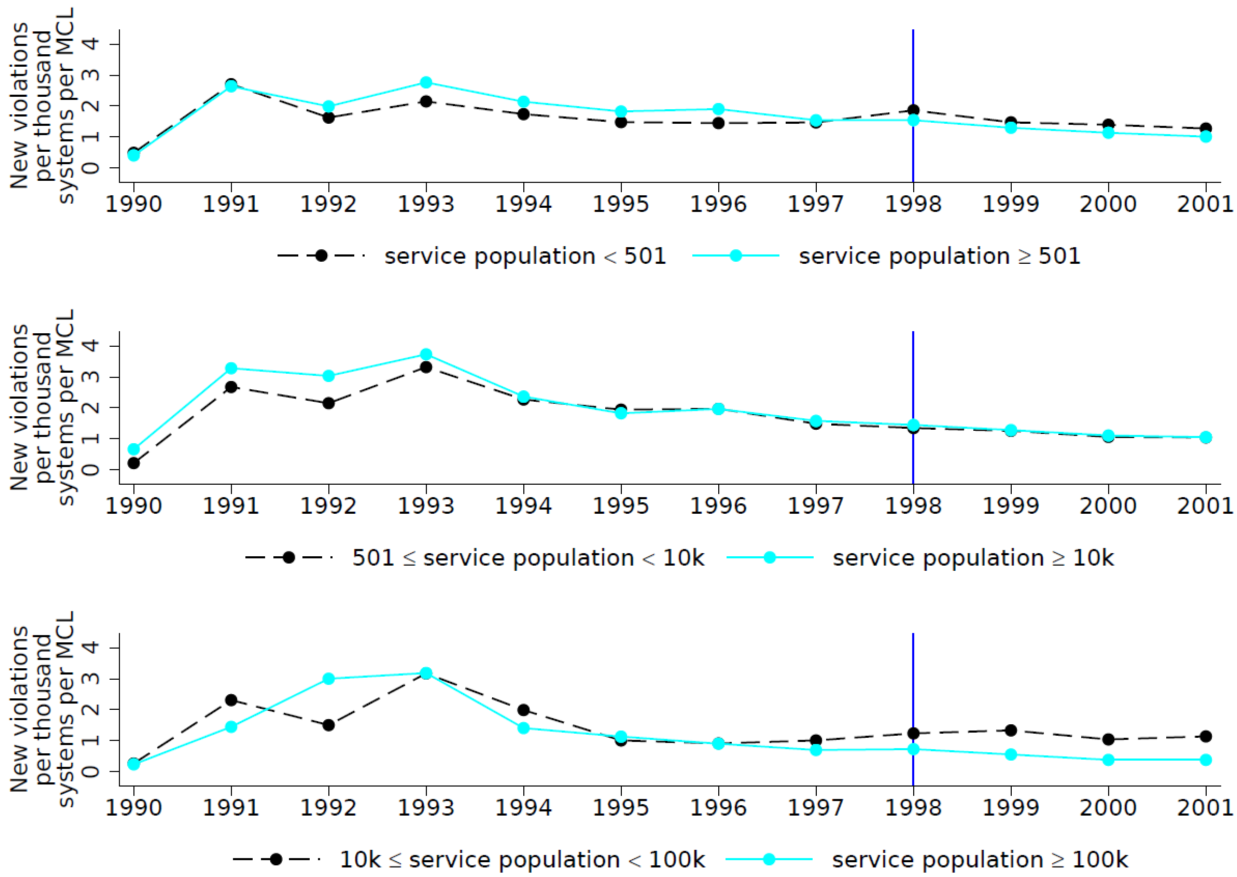

Caption: Figure 1 provides subtle evidence of declining health-based violations in the raw data beginning in 19987 for the publishing (top) and web (bottom) thresholds, relative to systems not required to disclose in these manners.

Notes: Number of health-based water quality violations in each year of the panel for (top panel) community water systems serving less than 500 customers compared to those serving greater than 500 customers; (middle panel) systems serving less than 10,000 customers (but greater than 500 customers) compared to systems serving greater than 10,000 customers; and (bottom panel) systems serving less than 100,000 customers (but greater than 10,000 customers) compared to systems serving greater than 100,000 customers. We normalize violations by the total number of water systems within each population service category and the annual number of MCL rules (the number of MCL rules grew from 31 in 1990 to 83 in 2001).

Figure 2: Annual impact of the publishing, mailing, and online disclosure requirements for each year post-policy

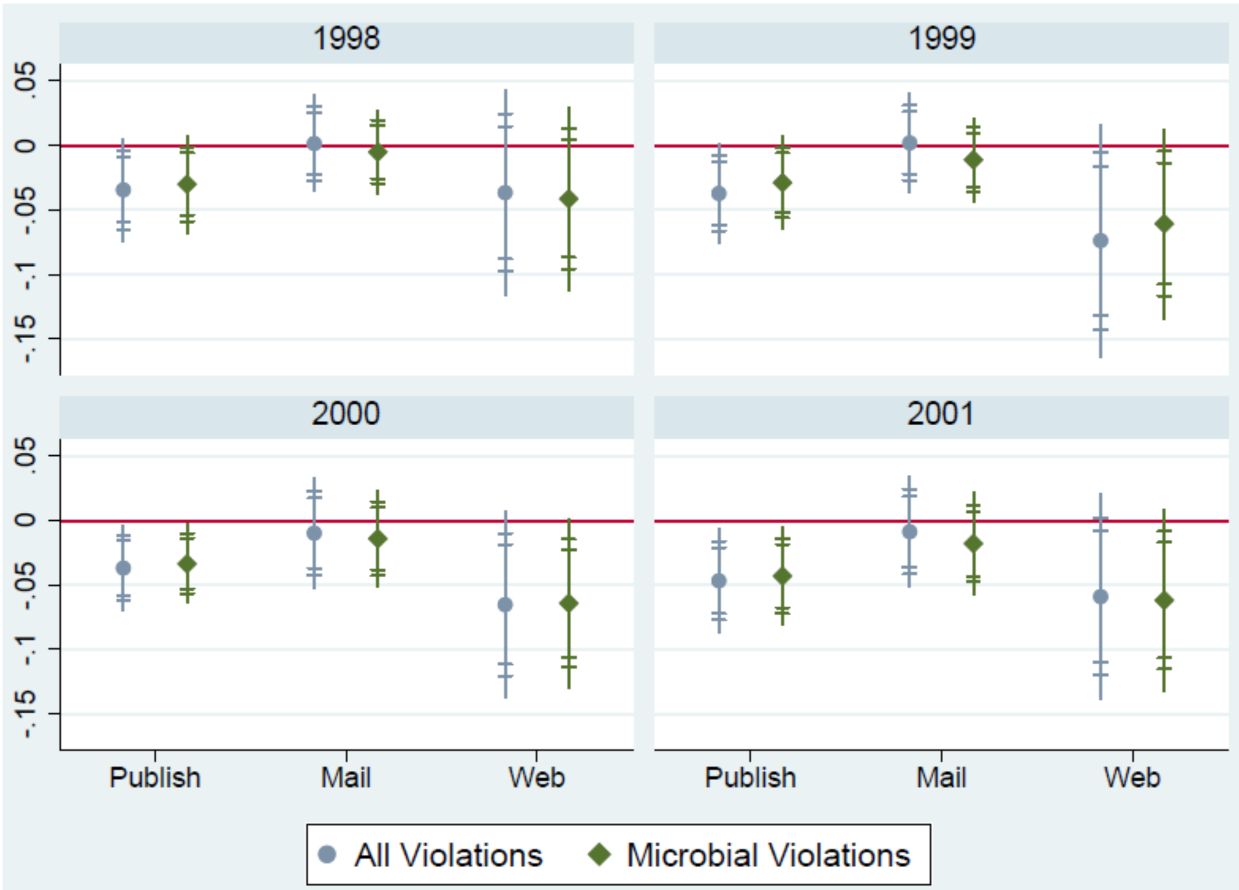

Notes: We graph point estimates from estimating Eq. (2) for each disclosure requirement separately, considering all health violations (gray-blue circle) and all microbial violations (green diamond) as the dependent variables. The 90th, 95th, and 99th percentile confidence intervals are marked with the inner hash mark, outer hash mark, and line (respectively).