

Assessing the Impact of PFAS Water Regulation

Laura Alcocer Quiñones *

University of California, Davis

October 2025

([Click here for latest version](#))

Abstract

Policy evaluation for emerging pollutants is challenging because monitoring is limited and data are sparse and censored. This paper demonstrates that failure to account for interval censoring can lead to misguided policy conclusions and proposes a method to overcome this challenge. I study this in the context of stricter state regulations for PFAS, a suite of emerging harmful water pollutants. A main obstacle is that water quality tests are interval-censored, so we only observe bounds on the distribution of PFAS concentration. I first show that without further assumptions, treatment effects are not generally point-identified and applying standard approaches that ignore this issue can lead to contradictory policy conclusions. I overcome this challenge using a parametric approximation to recover the latent cumulative distribution function through censored maximum likelihood. I implement the changes-in-changes approach using the recovered distribution of concentration to estimate the impact of tightening regulatory standards on water quality. For existing facilities, I cannot conclude that notification level changes had an impact on water quality across the distribution. I find suggestive evidence of investment in the form of new and retrofitted treatment plants to address PFAS after regulations became more stringent. The methods implemented allow for policy evaluation under interval censoring in other settings in economics.

*PhD Candidate in the Department of Agricultural and Resource Economics at the University of California, Davis (email: lalcocer@ucdavis.edu).

I am grateful to my committee, Dalia Ghanem, Katrina Jessoe, and Kevin Novan, for their continued support of this project. I thank Wes Austin, Jamie Hansen-Lewis, Eric Edwards, and seminar participants at the Allied Social Science Associations annual meeting, the Association of Environmental and Resource Economists summer conference, the Social Cost of Water Pollution Workshop, and the Agricultural and Applied Economics Association annual meeting for their comments. I thank Laleh Rastegarzadeh and Wendy Linck from the California State Water Resources Control Board for their help with institutional details.

1 Introduction

Measuring and monitoring new contaminants in drinking water is difficult due to technological limitations and heterogeneous data collection efforts. These constraints make policy evaluation for these types of pollutants challenging. Our inability to accurately detect concentrations at low levels leads to large amounts of censored data, which results in the loss of observations on part of the distribution of these contaminants. Data censoring poses a prominent challenge in the design of Per- and Polyfluoroalkyl Substances (PFAS) regulation.

PFAS, also known as “Forever Chemicals”, are a class of contaminants that are ubiquitous, long-lasting and difficult and costly to remove. These chemicals are known to pollute water supplies used for livestock, crop production and drinking water. Prolonged exposure has adverse consequences on human health and economic outcomes, including worsening fertility and infant health outcomes (Waterfield et al., 2020), kidney, breast, and testicular cancer (National Academies of Sciences and Medicine, 2022), and reduced long-run earnings from firefighting foam leaching (Jacqz, Somuncu, and Voorheis, 2024). Drinking water remains one of the primary exposure pathways (Domingo and Nadal, 2019). The Safe Drinking Water Act (SDWA) will bring two chemicals in this class under its purview starting in 2031. In response to growing concerns, states have also chosen to regulate PFAS. The impact that these regulations have on drinking water quality and human health remain uncertain.

To begin to understand the links between regulatory stringency and water quality, this paper takes advantage of state-level PFAS policies to evaluate the effect of tightening regulation on drinking water quality. I study changes in regulatory stringency in California and New Jersey, two states that implemented PFAS policies. I compare a tightening of regulation in California with no change in regulation in New Jersey. I analyze the effects of this policy change on the intensive and extensive margins. First, I assess the impact on drinking water quality in existing wells and treatment plants. Then, I examine whether water systems add and retire wells or treatment plants in response to the policy.

A main complication to the evaluation of these policies is data censoring from reporting limits. Reporting limits determine the minimum threshold level where PFAS can be measured with confidence and are set by laboratories or states. If a test measures concentration below the limit, the water system does not need to report the concentration. The implication of this is that if a test is interval-censored, I only observe that it occurred but do not observe the pollution reading. Censoring is pervasive in PFAS with approximately 60% of observations

interval-censored. When data are interval-censored, we can only bound the distribution of PFAS concentration and cannot point-identify treatment effects in general.

I demonstrate that in this setting, failure to account for interval censoring can lead to contradictory policy conclusions. How the econometrician treats censored data in estimation is the source of these contradictions. Using a difference-in-differences design, if I assume that censored observations are equal to zero, I find that, on average, the policy had no impact on water quality. However, if I assume that censored observations are equal to the reporting limit, I find that, on average, the policy significantly reduced PFAS concentration. Another common suggestion when dealing with interval-censored data is to drop censored observations for estimation. Doing so leads to the conclusion that the policy had a moderate impact in reducing pollutant concentration in water. Using a changes-in-changes approach (Athey and Imbens, 2006) to estimate the effects at the average under the same three scenarios also yields different estimates, in both magnitude and significance. These findings highlight the need to appropriately address interval censoring in policy evaluation.

To overcome the challenge posed by data censoring, I propose an approach that pairs censored maximum likelihood estimation with changes-in-changes. First, I fit the uncensored outcome to a Weibull distribution. Then, I use the recovered cumulative distribution functions in changes-in-changes to estimate the counterfactual distribution of concentration in the absence of a notification level change. I then estimate quantile treatment effects on the treated from stricter notification levels for PFOA and PFOS in California. I apply the methods described to a panel at the facility-quarter level for PFOA and PFOS drinking water samples in California and New Jersey from 2018 to 2020.

I find that, for existing facilities, I cannot provide evidence that the notification level change led to a reduction in PFAS concentration. This is due to the fact that California had a high degree of censoring in the pre-period and the previous regulatory threshold was located in the part of the distribution we do not observe, making it difficult to detect impacts in the lower part of the distribution, where we would expect them to occur. This result has important implications for how we set policies that align with technological advances to allow for accurate monitoring and compliance.

Water systems may also respond to regulation by adding or retiring wells and treatment plants. I construct a panel of the universe of wells and treatment plants testing for PFAS, including additions and retirements, in California from 2018 through 2022. Descriptive ev-

idence suggests that new plants and new wells are being constructed and new plants are positively correlated with the tightening of PFAS regulation. Using county-level data, preliminary insights show that new treatment plants tend to be in urban, more densely populated and higher-than-average household income areas. New wells tend to be cleaner, on average, relative to existing wells. This finding suggests water systems may be careful about placement of new wells regarding contamination.

The methodological contribution of this paper applies to a myriad of settings. Past literature on methods to address interval censoring include parametric and semiparametric methods for estimation and inference (Manski and Tamer, 2002; Abrevaya and Muris, 2020; Chay and Powell, 2001; Cook and McDonald, 2013) and imputation (Horowitz and Manski, 1998; Jenkins et al., 2011; Rios-Avila, Canavire-Bacarreza, and Sacco-Capurro, 2024). The literature has proposed methods that recover or bound treatment effects, both at the average and throughout the distribution. I contribute by pairing a parametric approximation with quasi-experimental methods to estimate a different object: the counterfactual distribution. A differentiating factor is that the approach proposed does not impute values for concentration but rather approximates the distribution under parametric assumptions and allows researchers to use standard methods for causal inference (e.g., difference-in-differences). The approach proposed in this paper can be applied to other settings in economics.

My paper also contributes to the literature on policy and regulatory design for water quality. A broad body of research on this topic includes the impact of loans and grants on water pollution (David A Keiser and Shapiro, 2019; D. Keiser et al., 2023), the impacts from mandated provision of information (Baker, L. Benneer, and S. Olmstead, 2023; L. S. Benneer and S. M. Olmstead, 2008), the impact of improvements to treatment infrastructure (Cutler and Miller, 2005), the implications of decentralizing regulation (David A Keiser, S. M. Olmstead, et al., 2021), and the consequences of regulatory failures (Christensen, David A. Keiser, and Lade, 2023). While past work has examined manipulation (L. S. Benneer, Jessoe, and S. M. Olmstead, 2009; Andarge et al., 2025) and measurement error (D. Keiser, 2019) in water quality data, there is limited evidence on the role of interval censoring on policy design and evaluation. I address this gap by providing evidence on how interval censoring in water quality measurements has implications for setting regulatory thresholds that can be accurately measured and monitored to ensure compliance. I demonstrate that failure to account for interval censored data can lead to contradictory policy conclusions.

In addition, this study assesses the impacts of regulatory stringency for PFOA and PFOS,

two understudied and newly-regulated pollutants. I contribute to the literature on the impacts of drinking water quality regulation which has focused on well-known pollutants such as nitrates (Hadachek, 2025; Pennino, Compton, and Leibowitz, 2017), coliform bacteria (Marcus, 2022; Marcus, 2025), arsenic (Foster et al., 2019), and phosphorus (Meyer and Raff, 2024).

The paper proceeds as follows. First, I discuss background on PFAS in drinking water and the context surrounding the regulation of these contaminants in California and New Jersey (Section 2). This section also describes the data used and its associated challenges. Next, I analyze the impact of this policy on the intensive margin, focusing on the implications of interval censoring and using a novel approach to overcome this issue (Section 3). This section also describes the identification strategy, its implementation and results for the intensive margin. Section 4 describes the analysis for the extensive margin and the implications of this policy change on infrastructure investments and their relationship with socioeconomic and demographic characteristics. Section 5 concludes.

2 Background and Data

2.1 Institutional background

PFAS are a family of man-made chemicals invented in the 1940’s for industrial purposes. A primary exposure pathway for humans is drinking water consumption (Sunderland et al., 2019; DeLuca et al., 2022). Exposure has been associated with certain cancers and birth deformities (National Academies of Sciences and Medicine, 2022). These substances can enter the water through landfill runoff, wastewater discharge, and firefighting foam runoff, among others. PFAS removal from water is often costly and difficult due to their longevity and persistence. Despite this, PFAS concentration can be reduced through dilution or treatments (Vo et al., 2020).

Public water suppliers are tasked with monitoring and sampling water. For most pollutants, including PFAS, samples come from a facility as defined by the Safe Drinking Water Information System (SDWIS) database (Environmental Protection Agency, 2018-2022). These facilities include wells, surface water intakes, treatment plants, distribution systems, storage tanks, etc. Water suppliers may sample at multiple points within a facility, for example, between leads and lags for treatment and at the entry point of the distribution system. When and where these samples take place depends on the regulation, and for the majority of pol-

lutants this is determined by the SDWA. However, federal regulation for PFOA and PFOS is not enforceable until 2031, allowing states to make these determinations in the meantime.

In the absence of enforceable federal regulation, states like California and New Jersey chose to regulate PFAS. In 2018, the California Department of Drinking Water (CA DDW) implemented a two-part policy using notification and response levels for two pollutants in this class of chemicals, PFOA and PFOS. While notification and response levels are not formally recognized as regulation, they are applicable to all water samples across the state and have actionable consequences when exceeded. Notification levels are thresholds that, when exceeded, require informing the utility’s governing body as well as their customer’s local government of the presence of PFAS in their supply. It does not require further action aside from providing information on detection and it does not require public notification. Response levels are thresholds that, when exceeded, require utilities to inform the public, retire that facility, or invest in treatment to lower PFAS concentration. The combination of notification and response levels is similar to setting a maximum contaminant level (MCL), the standard instrument to regulate water using a threshold for concentration.

This paper focuses on more stringent notification levels for PFOA and PFOS. In August 2019, the CA DDW tightened notification levels for these chemicals. California was not the only state to have enforceable drinking water standards for PFAS; New Jersey was among states leading the charge. New Jersey had informal regulation in the form of drinking water guidance. This threshold, much like California’s notification level, did not require action from water utilities when exceeded. Because of this similarity, New Jersey is comparable to California. However, New Jersey passed MCLs in June 2020, no longer making this comparison valid after that time period. Table 1 shows the regulatory changes for the contaminants studied in both states. These changes determine the sample that is used for estimating impacts for existing facilities.

A challenge that arises from using water quality data is the presence of reporting limits. These are thresholds set by laboratories or regulators that determine whether concentration can be measured with confidence. If a sample taken within a public water system is below the reporting limit, the water system reports that a test took place but not necessarily the concentration measured. Observations that are censored have a concentration that falls between zero and the reporting limit. For example, if a test for a well yields a concentration of 3 parts per trillion (ppt) and the reporting limit is 5 ppt, then the water utility does not need to report that concentration. If they do not report it, the data point shows a test was

Table 1: Regulatory Changes Related to PFAS in California and New Jersey

Date of Change	State	Contaminant	Type of Instrument	New Level	Initial Level
August 2019	CA	PFOA	Notification Level	5.1 ppt	14 ppt
August 2019	CA	PFOS	Notification Level	6.5 ppt	13 ppt
February 2020	CA	PFOA	Response Level	10 ppt	70 ppt
February 2020	CA	PFOS	Response Level	40 ppt	70 ppt
June 2020	NJ	PFOA	MCL	14 ppt	40 ppt
June 2020	NJ	PFOS	MCL	13 ppt	40 ppt

Notes: For California, a notification level exceedance can occur with any test performed at any sampling point. A response level exceedance is determined using the Quarterly Running Annual Average (QRAA). This requires that PWS average sample results during four calendar quarters. New Jersey’s initial level is drinking water guidance that eventually was reduced and replaced by an MCL.

taken on that date with the concentration missing, but the reporting limit for the test is reported. If they do report it, the data point shows the test taken and the concentration for that test. As technological advances for detection are made, reporting limits become smaller. Our ability to detect lower concentrations of PFAS has greatly increased over time with dramatic reductions in reporting limits relative to as early as five years ago. Before the policy change, California had reporting limits of up to 20 and 40 ppt for PFOA and PFOS, respectively. At that time, the notification levels for PFOA and PFOS were set at 14 and 13 ppt, respectively. This means that for some readings in the state, the reporting limit was higher than the regulatory threshold in place, making it impossible to identify if an exceedance had occurred.

2.2 Data

The primary data are composed of drinking water quality readings at wells and treatment plants. I supplement this with census data on socioeconomic and demographic characteristics and quarterly facility-specific characteristics. I construct a panel of drinking water samples for PFOA and PFOS for 2,173 facilities between 2018 and 2020 for California and New Jersey.

Drinking Water Quality: I use quarterly average PFOA and PFOS concentration at facilities as my outcome of interest.¹ This data was publicly accessed through the California State Water Resources Control Board (2018-2022) and provided through an Open Public Records Act request to the New Jersey Department of Environmental Protection.² Reporting limits

¹For facilities that have multiple sampling points (i.e., treatment plants), I also create measures of maximum and minimum PFAS concentration over all sampling points for robustness. The concern with using the average is that treatment plants test for PFAS in points other than the entrance to the distribution system and this statistic may be insufficient to characterize how to weight other sampling points in estimation.

²While the ideal data would be measuring PFAS concentration experienced by consumers at the tap, the

lead to large amounts of interval censoring in this setting. Approximately 60% of observations for both states are interval-censored between zero and the reporting limit. For the tests that we do observe, approximately 30% of them would be considered violations under the SDWA for PFOA and PFOS.

Due to interval censoring, I create two new variables to define the bounds of concentration. The lower bound on the concentration is created by replacing censored observations with zeros. The upper bound is created by replacing censored observations with the reporting limit. Summary statistics for the bounds of the concentration and the reporting limits over the sample period are provided in Table 2. Fixing a pollutant, the average concentration when censored observations are replaced with zeros is different than the average concentration when censored observations are replaced with the reporting limit. This difference is particularly stark for California. On average, reporting limits for both pollutants are higher in California.³

Table 2: Summary Statistics: 2018-2020

	California			New Jersey		
	Mean	SD	N	Mean	SD	N
Cens. PFOA Replaced w/Zero (ppt)	7.41	18.63	3,358	4.17	15.18	6,697
Cens. PFOA Replaced w/RL (ppt)	9.85	18.27	3,358	5.19	14.95	6,697
PFOA Reporting Limit (ppt)	4.36	5.61	1,827	1.77	1.11	3,872
Cens. PFOS Replaced w/Zero (ppt)	13.91	35.97	3,358	4.12	25.44	6,693
Cens. PFOS Replaced w/RL (ppt)	17.53	35.85	3,358	5.27	25.28	6,693
PFOS Reporting Limit (ppt)	7.17	12.48	1,669	1.74	1.12	4,409

Notes: The unit of observations is facility-quarter. Summary statistics are provided for the bounds of the distribution and reporting limits for California and New Jersey from 2018 through the first quarter of 2020. Columns shows the mean, standard deviation, and count of non-missing observations for each group. Reporting limits are provided for readings with censored concentration.

A second data issue is that water quality data suffers from large amounts of attrition. Because public water systems are not required to test all their facilities consistently, I observe many missing values. In this setting, attrition can occur for two reasons: cleaner facilities test less, and dirty facilities can be taken offline.⁴ Table A2 shows the count of California regulation and testing occurs at facilities within public water systems.

³Several readings had reporting limits as high as 20 and 40 ppt for PFOA and PFOS, respectively.

⁴Per California regulation, if a facility has four consecutive quarters with non-detections, then it can submit a request to change its monitoring schedule. There are approximately 500 facilities in California that qualify for a change in monitoring schedule based on available data. This constitutes approximately 55% of

facilities by activity status and monitoring schedule change eligibility. Since the majority of facilities that test for these pollutants are not eligible for the schedule change nor are they inactive, there may be concerns that attrition is likely non-random. To address this issue, I subset my sample to a balanced panel of facilities that test at least once before and after the policy change.

Utility Characteristics: I use a combination of data provided by the California State Water Resources Board (2025) and the federal government from the SDWIS database. I obtain time-varying characteristics for public water system facilities, including when a facility was active, if and when it was taken offline.

Socioeconomic and Demographic Characteristics: I use annual data at the county level from the American Community Survey (U.S. Census Bureau, 2018-2022).⁵ Characteristics include percentage of non-white population, average household income, number of households, and housing units, among others. Additionally, I use the Census Bureau’s urban-rural classification (U.S. Census Bureau, 2020) to determine which water systems are in these areas.

3 Intensive Margin

First, I assess whether facilities operating before and after the policy change are being cleaned up by water systems in response to a notification level change. This section describes the practical implications of using interval-censored data and that ignoring or failing to address this issue in a critical way can result in incorrect policy conclusions. It then proposes an approach to applying changes-in-changes in the presence of interval censoring and shows the results from a change in notification levels. For ease of exposition, the analysis shown in the main body of the paper corresponds to PFOS. Analysis for PFOA can be found in Appendix A.

observed facilities in the state. The second potential reason for attrition relates to a facility’s status. If a facility is tested and it is found in violation of the notification level or response level, then the corresponding water system may choose to shut it down temporarily or permanently depending on their ability to treat it. If a facility is inactive, it is not required to test but may choose to do so to return to active status.

⁵Because of the COVID-19 pandemic, the year 2020 was not provided and I use 2019 estimates for that year.

3.1 Practical implications of interval censoring

3.1.1 Implication I: We only observe bounds on PFAS distribution

Let Y_{it}^* denote the concentration of PFAS in drinking water for facility i at time t , $S_{it} = 1$ if a public water system (or laboratory) chooses not to report the concentration for facility i and define r_t as the reporting limit at time t . We observe Y_{it} which is a partially interval-censored version of Y_{it}^* , specifically

$$Y_{it} = \begin{cases} Y_{it}^*, & \text{if } S_{it} = 0 \text{ or } Y_{it}^* > r_t \\ [0, r_t], & \text{if } S_{it} = 1 \text{ \& } Y_{it}^* \leq r_t \end{cases}$$

An observation is censored if concentration falls below the reporting limit and the public water system (or laboratory) chooses not to report it.

Due to interval censoring, we can only bound the distribution of PFAS concentration using the observed data. Let $F_{Y_{it}^r}(y)$ denote the distribution of PFAS concentration when censored observations are replaced with the reporting limit, and $F_{Y_{it}^z}(y)$ denote the distribution of PFAS concentration when censored observations are replaced with zeros, where

$$Y_{it}^r = \begin{cases} Y_{it}^*, & \text{if } S_{it} = 0 \text{ or } Y_{it}^* > r_t \\ r_t, & \text{if } S_{it} = 1 \text{ \& } Y_{it}^* \leq r_t \end{cases}$$

$$Y_{it}^z = \begin{cases} Y_{it}^*, & \text{if } S_{it} = 0 \text{ or } Y_{it}^* > r_t \\ 0, & \text{if } S_{it} = 1 \text{ \& } Y_{it}^* \leq r_t. \end{cases}$$

We do not observe the latent distribution of PFAS concentration but rather its bounds. The distribution of concentration ($F_{Y_{it}^*}(y)$) falls between $F_{Y_{it}^z}(y)$ and $F_{Y_{it}^r}(y)$. Formally,

$$F_{Y_{it}^*}(y) \in [F_{Y_{it}^r}(y), F_{Y_{it}^z}(y)] \text{ for } y \leq r_t^{max}$$


$$F_{Y_{it}^*}(y) = F_{Y_{it}^r}(y) = F_{Y_{it}^z}(y) \text{ for } y > r_t^{max}$$

where r_t^{max} is the maximum reporting limit in period t for $t \in \{0, 1\}$.

To illustrate this implication, Figure 1 shows the bounds for the distribution of PFOS concentration in California before and after the policy change. The blue line corresponds to the case when censored observations are replaced with the reporting limit, the red line corresponds to the case when censored observations are replaced with zeros, and the dotted lines

correspond to the notification level before and after the change. The bounds converge past the maximum reporting limit for the period, since censoring occurs at the lower tail of the distribution. While we do not observe the uncensored distribution, we observe its bounds. What is interesting in this figure is where notification levels are set. Before the policy change, the threshold was in the unobserved portion of the distribution (i.e. the censored portion). After the change, the threshold was in the observed portion. This has important implications for how we set policy standards that can be appropriately measured and enforced, especially in the context of emerging contaminants.

Figure 1: Bounds for PFOS Concentration in California Before and After Policy Change



Notes: Each panel shows the bounds for the distribution of PFOS concentration in California before and after the policy change, respectively. Lower bound corresponds to the case where censored observations are replaced with the reporting limit and the upper bound corresponds to the case where censored observations are replaced with zeros. The dashed vertical line indicates the notification level in place at that time.

3.1.2 Implication II: Treatment effects are not point-identified in general

From implication I, we only observe the upper and lower bounds on PFAS concentration. As a result, treatment effects are not point-identified in general. To provide a specific example,

let us consider the average treatment effect on the treated (ATT). We know

$$ATT = E[Y(1)|D = 1] - E[Y(0)|D = 1] \quad (1)$$

where $D = 1$ is an indicator for whether a facility is treated, i.e., located in California. We know $E[Y(1)|D = 1] \in [\mu_1^r, \mu_1^z]$ and $E[Y(0)|D = 1] \in [\mu_0^r, \mu_0^z]$ where μ_0^j is mean concentration for the untreated conditional on being treated, μ_1^j is mean concentration for the treated conditional on being treated, and $j \in \{r, z\}$ with r denoting the lower bound and z denoting the upper bound. From equation 1,

$$ATT \in [\mu_1^r - \mu_0^z, \mu_1^z - \mu_0^r] \quad (2)$$

Without further assumptions, we can only bound the ATT.

3.1.3 Implication III: Naïve counterfactual analysis may lead to incorrect conclusions

To implement existing approaches, such as difference-in-differences (DID) or changes-in-changes (CIC) introduced in Athey and Imbens (2006), one must observe the outcome distribution, and not solely the bounds on it. To operationalize them, one might wonder what DID or CIC would yield if they were applied to $Y_{it}^z, Y_{it}^r, Y_{it}^*|S_{it} = 0$, or other commonly used cases including replacing censored observations with half of the reporting limit ($Y_{it}^{\frac{r}{2}}$) and the square root of half of the reporting limit ($Y_{it}^{\sqrt{\frac{r}{2}}}$). Unbiased estimates require assuming that parallel trends hold for each of these cases.⁶ Suppose the required identifying assumptions hold for the cases described. The estimating equation for a panel of the concentration at facility i in quarter-year t is:

$$Y_{it} = \alpha_0 + \beta_1 CA_i + \beta_2 Post_t + \beta_3 CA_i Post_t + \varepsilon_{it} \quad (3)$$

where the outcome Y_{it} includes the definitions of PFOS concentration described, CA_i is an indicator for whether facility i is in California and $Post_t$ is an indicator for whether quarter-year t is in the post period. The coefficient of interest is β_3 , measuring the effect of a change in California's notification levels for PFOA and PFOS in August 2019. New Jersey is used as a control group.

Table 3 reports the results from this naïve model using DID (row 1) and CIC (row 2). The

⁶Figure A9 shows these trends for the cases studied.

difference across columns consists of how censored data is treated. For DID, the results are mixed. When using the upper bound of the distribution, the results suggest the policy had no impact on water quality. However, when using the lower bound of the distribution, the results suggest the policy led to a significant improvement in water quality. Replacing censored observations with half of the reporting limit also suggest improvements but replacing censored observations with the square root of half of the reporting limit yields no impact. When using CIC, the results change in magnitude, sign and significance. This finding highlights how the treatment of censored observations affects policy evaluation and can lead to contradictory policy conclusions.

Table 3: Average Treatment Effects for PFOS Concentration Using DID and CIC

	(1)	(2)	(3)	(4)	(5)
	Y_{it}^z	Y_{it}^r	$Y_{it}^* S_{it} = 0$	$Y_{it}^{\frac{r}{2}}$	$Y_{it}^{\sqrt{\frac{r}{2}}}$
ATT DID	-1.365 (1.073)	-6.309** (2.070)	-5.626* (2.395)	-3.807** (1.436)	-1.852 (1.105)
ATT CIC	0.551 (1.786)	-3.680 (2.279)	-0.787 (3.568)	-1.360 (1.818)	0.305 (1.712)
Observations	10051	10051	3973	10051	10051
N-PWS	1224	1224	554	1224	1224

Notes: The unit of observations is facility-quarter-year. The sample includes facilities within PWS testing for PFAS in California and New Jersey between 2018 and the first quarter of 2020. The independent variable is an indicator for the post period. The dependent variables are the upper bound of PFOS concentration (column 1), the lower bound of PFOS concentration (column 2), partially interval-censored PFOS concentration (column 3), censored PFOS concentration replaced with half of the reporting limit (column 4) and censored PFOS concentration replaced with the square root of half the reporting limit (column 5). CIC estimates bootstrapped with 1,000 replications. Standard errors clustered at the PWS level are reported in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Furthermore, a shortcoming of implementing the approach described is that these estimates do not bound the ATT. Suppose we take $F_{Y_{it}^z(y)}$ as the uncensored distribution of PFAS concentration. Then the treatment effect is $\mu_1^z - \mu_0^z$. Now suppose we take $F_{Y_{it}^r(y)}$ as the uncensored distribution of PFAS concentration. Now the treatment effect is $\mu_1^r - \mu_0^r$. These bounds fall within those derived in implication 2. This fact implies that the estimates in column 1 do not provide an upper bound for the ATT and analogously, the estimates in column 2 do not provide a lower bound for the ATT.


Threshold policies, like the one evaluated, tend to induce different responses above and below the threshold. The policy will likely only affect a part of the distribution and commonly used

methods, such as DID, only capture these impacts at the average. Using CIC, I estimate the counterfactual distribution and the corresponding quantile treatment effects of the policy.⁷ Quantile treatment effects on the treated (QTT) are defined as the vertical difference between the observed and counterfactual concentration after fixing a quantile.

Figure 2 shows the QTT for PFOS concentration. The results are sensitive to the definition of PFOS concentration used for point-identification. The first panel corresponds to Y_{it}^z . For the upper bound, the policy had largely no impact on water quality with only one quantile showing increased concentration. The second panel corresponds to Y_{it}^r . For the lower bound, the policy moderately decreases water quality for places with lower amounts of PFAS but has large effects in the expected direction for dirtier facilities. The third panel corresponds to $Y_{it}^*|S_{it} = 0$. Although the distribution of quantile treatment effects reflects larger decreases in concentration for dirtier locations, a large proportion of the quantiles show no impact of the policy on water quality. The last two cases are commonly used approaches to dealing with interval-censored data. The panel corresponding to $Y_{it}^{\frac{r}{2}}$ shows an increase in water quality around the 80th quantile while the panel corresponding to $Y_{it}^{\sqrt{\frac{r}{2}}}$ shows a flat insignificant response throughout the distribution. I obtain qualitatively different results depending on the outcome used. Using naïve counterfactual analysis, as shown in this section, can lead to misleading policy conclusions. This finding emphasizes the need to appropriately address interval censoring in policy evaluation.

⁷The counterfactual and observed distributions for PFOS concentration are shown in Figure A11 for the three primary cases.

Figure 2: Quantile Treatment Effects on the Treated with Changes-in-Changes for PFOS



Notes: Each panel shows the quantile treatment effects on the treated for PFOS concentration from the 10th to the 90th quantile. Top panel on the left corresponds to the case where censored observations are replaced with zeros (upper bound), the top middle corresponds to the case where censored observations are replaced with the reporting limit (lower bound), the top right corresponds to the case where censored observations are dropped, the bottom left corresponds to the case where the censored observations are replaced with half of the reporting limit and the bottom middle corresponds to the case when censored observations are replaced with the square root of half the reporting limit. The last two cases are common suggestions among practitioners.

3.2 Applying CIC in the presence of interval censoring

DID estimates the average treatment effect on the treated but fails to provide information on the effects of the policy for the rest of the distribution. Policy standards that rely on thresholds may induce different responses throughout the outcome's distribution. In this setting, there is interest in understanding the effect of this change on concentrations for facilities with lower amounts of PFAS versus facilities with large traces of these substances in their water. Furthermore, regulation may induce a large proportion of facilities to bunch just underneath the new threshold. These effects are muddled at the average, making DID flawed in this context. Instead, I use CIC to estimate the counterfactual distribution of PFAS concentration in the absence of the regulatory change, and consequently, quantile treatment effects. The main challenge to address when implementing this method is the presence of a partially interval-censored outcome such as PFAS concentration in drinking water.

To overcome this issue and estimate the effect of this policy, I pair a parametric approach with CIC. This parametric approximation fits a Weibull distribution to the observable concentration data to recover the latent cumulative distribution function (CDF). Using censored maximum likelihood (MLE), I estimate the CDF for each period and each group based on Ghanem, Shen, and Zhang (2020). These CDFs are used in CIC to quantify the impact of tightening notification levels on water quality.

3.2.1 Identification strategy

Let Y_{it}^* be PFAS concentration in facility i and time t . Each facility i belongs to group $G_i \in \{0, 1\}$, where $G_i = 1$ denotes the treatment group and $G_i = 0$ denotes the control group. Facilities are observed for a pre-treatment period and a post-treatment period. The observed outcome for facility i at time t can be written as a function of the potential outcomes:

$$Y_{it}^* = G_i Y_{it}^*(1) + (1 - G_i) Y_{it}^*(0)$$

where $Y_{it}^*(1)$ is the potential outcome under a regulatory change and $Y_{it}^*(0)$ is the potential outcome in the absence of a regulatory change.

Assumption 1. CIC assumptions on latent outcome:

1. Water quality for a facility in the absence of a change in regulation satisfies the equation: $Y_{it}^*(0) = m_t(U_{it})$.
2. Monotonicity: Function $m_t(u)$ is strictly increasing in u for $t \in \{0, 1\}$.
3. The distribution of unobservables conditional on treatment is time invariant, $U_{i0}|G_i \stackrel{d}{=} U_{i1}|G_i$.
4. The support of $Y_{it}^*(0)|G_i = 1$ is included in the support of $Y_{it}^*(0)|G_i = 0$ for $t = 0, 1$.

The first assumption defines water quality for facility i as a function of unobservable characteristics (U_{it}). It requires that absent treatment, if unobservable characteristics for facilities i and j are the same ($U_{it} = U_{jt} = u$) then baseline water quality is the same for both facilities. The second assumption requires that higher water quality correspond to a higher value of u . If water quality is determined by effort to curb pollution in specific facilities, then the monotonicity assumption ensures water quality is increasing in effort. The third assumption states the population of treated and untreated facilities can have a different distribution of the unobservable but the distribution of the unobservable for each group is stable across

time. CIC conditions for continuous outcomes are equivalent to a scale-invariant version of parallel trends (Ghanem, Kédagni, and Mourifié, 2023). Lastly, the support condition is standard.

Assumption 2. The CDF of $Y_{it}^*|G_i = g$ is $H(y; \theta_{gt})$ where $t \in \{0, 1\}$, $g \in \{0, 1\}$ and $H(y; \theta_{gt})$ is a known function with density $h(y; \theta_{gt})$.

For simplicity, define $H_{gt}(y) = H(y; \theta_{gt})$. Then $H_{11}(y) = F_{Y_{it}^*(1)|G_i=1}(y)$. For the remainder of the paper, let $H_{11}^{CF}(y) = F_{Y_{it}^*(0)|G_i=1}(y)$. Based on the assumptions above, I can identify the counterfactual distribution from the following:

$$H_{11}^{CF}(y) = H_{10}(H_{00}^{-1}(H_{01}(y))) \quad (4)$$

where $H_{10}(y)$ is the distribution of concentration for California in the pre-period, $H_{00}^{-1}(q)$ is the inverse of the distribution of concentration for New Jersey in the pre-period, and $H_{01}(y)$ is the distribution of concentration for New Jersey in the post-period.

3.2.2 Implementation

Estimation of the latent distribution is performed by treatment group and time period, recovering the distribution for control and treatment before and after the policy took place. Define Y_{igt} as interval-censored concentration for facility i belonging to group g at time t and r_{gt}^{max} as the maximum reporting limit at time t for group g .⁸ Let Q_{igt} be the number of quarters observed for facility i in group g and time t , and S_{gt} be the number of facilities in group g at time t . $H(\cdot; \theta_{gt})$ is the latent distribution of concentration for a contaminant – PFOA or PFOS. The log-likelihood function is the following:


$$\sum_{i=1}^{S_{gt}} \sum_{t=1}^{Q_{igt}} \{ \mathbb{1}\{Y_{igt} > r_{gt}^{max}\} \log h(y_{igt}; \theta_{gt}) + \mathbb{1}\{Y_{igt} \leq r_{gt}^{max}\} \log(H(r_{gt}^{max}; \theta_{gt})) \}$$

Figure 3 illustrates the output of the censored maximum likelihood estimation for PFOS in New Jersey and California before and after the notification level change. The x-axis is PFOS distribution in parts per trillion (ppt) and the y-axis is the associated probability. The yellow section highlights the portion of the distribution that is censored for the maximum likelihood estimation. Each panel shows three lines: the upper bound of the latent distribution (Y_{it}^z), the lower bound of the latent distribution (Y_{it}^r), and the fitted distribution. For both treatment (California) and control (New Jersey), the observed part of the distribution exhibits a

⁸To avoid complex estimation with constrained maximum likelihood, I assign observations below r_{gt}^{max} as censored at that level. While it effectively reduces the information used for fitting the Weibull distribution, it allows straightforward inference via bootstrap.

good fit within the 95% confidence interval. I take these CDFs as the distribution for PFOS concentration for each group and each period and use them to evaluate the policy of interest.

Figure 3: Fitted PFOS Concentration Distribution for New Jersey and California in Pre and Post Periods




Notes: Panel (a) corresponds to New Jersey (control) where the left figure shows the parametric fit for the time period before notification levels changed in California and the right panel corresponds to the fit for the time period after. Analogously, panel (b) corresponds to California (treatment). Both groups and time periods are fit using a Weibull distribution. The lower bound (dashed line) denotes the distribution of PFOS concentration replacing censored observations with the reporting limit. The upper bound (solid line) denotes the distribution of PFOS concentration replacing censored observations with zero. The blue line corresponds to the Weibull fit for the uncensored distribution of PFOA concentration. 95% confidence intervals are provided for Weibull fit and the uncensored portion of the distribution.

3.2.3 Results

Figure 4 shows the quantile treatment effects on the treated for PFOS along the distribution. While these results do not provide evidence for the hypothesis that the notification level reduction led public water systems to clean the water, there is an important caveat to this conclusion. Due to the high degree of censoring in the pre-period in California, the confidence interval in the lower part of the distribution is too wide. Reporting limits can affect one's ability to detect impacts if we expect these to be primarily in the low part of the distribution. This fact paired with our inability to appropriately measure PFAS concentration around the regulatory threshold leads to noisy estimation. This finding highlights the need for setting standards aligned with technological advancements for detection.

Figure 4: Quantile Treatment Effects on the Treated for PFOS



Notes: Quantile treatment effects on the treated accounting for interval-censored data. Solid black line denotes the average treatment effect on the treated using the method proposed. Gray area denotes the 95% point-wise confidence intervals on the QTT bounds. Standard errors are computed using 1,000 bootstrap replications with resampling conducted at the water system level.

Since the regulation does not force utilities to act in case of exceedance, it is possible that a notification level change was non-binding. Further robustness checks are conducted including a larger sample size with the inclusion of more states and using minimum and maximum measurements of concentration at treatment plants. Using a larger sample size, Figure A4 reports the fit from a Weibull distribution to the data. Figure A5 shows the results from

applying CIC with interval-censored data including other states as controls. Additionally, Figure A6 shows the results from applying CIC with interval-censored data using the minimum or the maximum concentration within a treatment plant. None of these specifications change the primary result. Details on robustness checks are provided in Appendix B.


4 Extensive Margin

Analyzing the intensive margin alone may not capture the full effects of the policy. Water systems can respond to tighter regulatory standards by adding or shutting down wells and treatment plants. To capture this response, I perform descriptive analysis on facility changes in California following the regulatory changes studied. I assemble a comprehensive quarterly panel for all treatment plants and wells in California that test for PFOA and PFOS between 2018 and 2022. I merge these with county level data on socioeconomic and demographic characteristics to study the relationship between these variables and new facilities across the state.⁹. Wells and treatment plants can come online or offline following the policy. I identify which facilities enter and exit the sample from 2018 through 2022.

First, water systems may add new wells or take them offline in response to the policy. After the policy, sources that test above the notification or response level may be taken offline to reduce exposure to PFAS. In contrast, new sources may appear in response to high degrees of contamination. For example, if a water system has sources in exceedance of the response level, they may drill new wells while they address contamination in other sources. I analyze this by counting the number of new and retired wells across time. Figure 5 shows the flow of these two counts. There are more well additions than removals. Neither of these measures appear to be correlated with the change in regulation.

⁹Summary statistics on PFAS concentration and socioeconomic and demographic characteristics for the sample are provided in Table A3


Figure 5: Wells Added and Retired



Notes: Count of newly active and retired wells over time. Using data on the universe of wells that test for PFAS in California, a well is considered newly active if it was not in the sample in the three prior quarters. Wells are spot checked to ensure they are reporting for the first time. Because wells can remain inactive for long periods of time, I refer to these as newly active wells. A well is considered retired if there is no data on it for 3 consecutive quarters. Dashed vertical lines indicating the change in notification levels (NL) and response levels (RL).

Although new wells do not appear to be correlated with the policy change, one might wonder if they are cleaner than existing ones in terms of PFAS concentration. This is important because water systems may be strategic about where they place new wells as a margin of response. I visually assess this comparison by plotting the quarterly average of the bounds of PFAS concentration for new and existing wells. The upper bound was computed by taking the mean over each quarter of concentration when censored observations are replaced with zeros. The lower bound was computed by taking the mean over each quarter of concentration when censored observations are replaced with the reporting limit. Figure 6 shows the range of average concentration for wells over time. For both contaminants, new wells are much cleaner than wells reporting before and after the policy. Starting in 2019, the range of PFAS concentration for wells reporting before and after the policy dramatically drops. This is likely due to the technological changes that allow for concentrations to be detected at lower levels which lead to tighter bands over time. Water systems are constructing or bringing back online wells that are substantially cleaner. This result could suggest strategic behavior on the location of these new wells.

Figure 6: Bounds for Average Concentration in New Wells vs. Wells in Intensive Margin After Notification Level Change




Notes: Each panel corresponds to the quarterly average of the bounds of PFAS concentration in the post-period. The upper bound was computed by taking the mean over each quarter of concentration when censored observations are replaced with zeros. The lower bound was computed by taking the mean over each quarter of concentration when censored observations are replaced with the reporting limit. Red bands denote wells in the intensive margin panel while blue bands denote new wells identified through the extensive margin panel.

Another margin of response for water systems is through the construction or retrofitting of treatment plants. Due to the permanence of PFAS in the environment, water systems may choose to address tighter regulation by increasing their capacity for treatment. To show this margin, I count the number of new and retrofitted treatment plants in California over time. I classify treatment plants as new or retrofitted using federal and state records for these facilities. Details on this classification are provided in Appendix D.2. Figure 7 shows this count over time. The vertical lines indicate when the notification and response levels for PFOA and PFOS changed, respectively. Before notification levels changed, only a single new treatment plant that tested for PFAS appeared in the sample. There is a stark increase in these types of facilities in the middle of 2020. This is suggestive evidence that the two-part policy led water systems to invest in treatment. News articles support this hypothesis as public water systems found their wells to exceed response levels which led to taking those facilities offline and subsequently, pushed them into investing in treatment plants.¹⁰

¹⁰One such example is Yorba Linda cleaning up PFAS from their water following large scale contamination

Figure 7: New and Retrofitted Treatment Plants



Notes: Count of new and retrofitted treatment plants per period. See Appendix D.2 for details on assigning a treatment plant as new or retrofitted. The dashed vertical lines indicate the change in notification levels (NL) and response levels (RL) for PFOA and PFOS in California.

Regulation was likely to influence the influx of new and retrofitted treatment plants, raising questions on where these new facilities are located and what the composition of surrounding communities looks like. There is concern that PFAS may disproportionately affect disadvantaged communities (Liddie, Schaider, and Sunderland, 2023; Maruzzo et al., 2025; Mueller et al., 2024). Understanding the relationship between the location of new and retrofitted treatment plants and socioeconomic and demographic characteristics is important to determine whether action is being taken to protect those bearing the highest risk. To understand this relationship, I analyze the difference in means for various county-level characteristics between existing and new or retrofitted treatment plants. The unit of observation is facility-quarter-year. Table 4 shows that counties with new plants tend to have a higher number of households and higher-than-average household income. The population in these counties tends to be less white and less hispanic. I analyze these differences in county-level characteristics for newly active and existing wells. Table A4 reports these results and shows that newly active wells tend to be located in counties with higher hispanic and black population.

after the policy was implemented. Find one of the articles on this topic here.

Table 4: Socioeconomic and Demographic Characteristics for Counties with New vs. Existing Treatment Plants

	(1)	(2)	(3)
	New or Retrofitted Plants	Existing Plants	Difference (1)-(2)
Num. of Households (thousands)	2313.8 (1252.8)	1286.4 (1212.3)	1027.32*** (0.000)
Average HH Income (USD)	118012.9 (14096.6)	105102.9 (19502.6)	12909.96*** (0.000)
Total Population (thousands)	6721.8 (3534.9)	3850.0 (3502.5)	2871.79*** (0.000)
Population Served by PWS (thousands)	113.2 (122.0)	148.4 (165.8)	-35.23*** (0.000)
Num. of Housing Units (thousands)	1008.4 (1467.7)	975.1 (1216.7)	33.29 (0.764)
% White Population	35.8 (8.8)	47.3 (14.7)	-11.46*** (0.000)
% Black Population	6.3 (2.3)	6.3 (2.2)	0.07 (0.693)
% Hispanic Population	42.3 (7.9)	43.8 (11.1)	-1.53** (0.017)
% Urban Area	1.0 (0.2)	0.9 (0.3)	0.06*** (0.001)
Observations	190	1,683	1,873

Notes: The unit of observation is facility-quarter-year. Table shows difference in means for demographic and socioeconomic characteristics between counties with new (or retrofitted) and existing treatment plants. Column (1) corresponds to counties with PWS that had new or retrofitted treatment plants and column (2) corresponds to counties with PWS that had existing treatment plants. Data corresponds to annual ACS matched with the extensive margin panel using principal county served by a public water system. The corresponding county averages are weighted by the number of facility-quarter observations for water systems located within each county. Standard deviations are reported for the means and p-values are reported for the differences: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Addressing PFAS contamination by installing new treatment, rehabilitating contaminated wells or drilling new wells can be costly. The varying response from water systems is likely correlated with their financial situation since utilities with limited resources may not have the capacity to address contamination. This has become an issue at the forefront of PFAS policy and one that has had heterogeneous responses from states. California has provided financial assistance in the form of grants and loans for systems to implement treatment and to rehabilitate contaminated wells. The majority of these grants were disbursed after 2022, but five projects were funded that year. These projects were only awarded a maximum of \$5 million each with requests ranging from \$24 to \$178 million. To analyze whether these grants were provided to water systems with higher degrees of contamination, I examine the difference in means for the bounds PFOA and PFOS concentration for water systems that received a grant and those that did not. The unit of observation for this analysis is a facility-

quarter. Table 5 shows that, on average, the concentration is higher for water systems that received a grant relative to those that did not. This result provides suggestive evidence that grants were allocated to water systems with higher contamination.

Table 5: Water Quality for PWS that Received or Did Not Receive Grants

	(1)	(2)	(3)
	Grant received	No grant received	Difference (1)-(2)
Censored PFOA Replaced with Zero (ppt)	8.07 (7.82)	5.21 (16.65)	2.86*** (0.000)
Censored PFOA Replaced with RL (ppt)	8.58 (7.34)	6.66 (16.36)	1.92*** (0.000)
Censored PFOS Replaced with Zero (ppt)	17.21 (19.01)	9.41 (25.48)	7.80*** (0.000)
Censored PFOS Replaced with RL (ppt)	17.67 (18.61)	10.94 (25.26)	6.73*** (0.000)
Observations	866	13,534	14,400

Notes: The unit of observation is facility-quarter-year. Table shows differences in PFAS concentration for public water systems with grants and those without between 2018 and 2022. A water utility is classified as having received a grant throughout the sample if, during 2022, they received funding for a project to address PFAS. Column (1) corresponds to observations for public water systems that received grant funding during the sample and column (2) corresponds to observations for public water systems never received grant funding during the sample. Standard deviations are reported for the means and p-values are reported for the differences: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

5 Conclusion

In this paper I estimate the impact of tightening regulatory standards for PFAS on water quality. I show that under interval censoring, we do not observe the distribution of PFAS concentration in water, treatment effects are not point-identified without further assumptions, and using interval-censored data in naïve counterfactual analysis can lead to misleading results. I overcome this challenge by pairing a parametric approximation with changes-in-changes to estimate the counterfactual distribution and treatment effects. Using the proposed approach, I find that on the intensive margin I am unable to detect that stricter standards led to changes in water quality. This is due to high degrees of censoring in the pre-period for California and the fact that the previous regulatory threshold was located in the censored portion of the distribution. This finding has strong implications for how we set policies that align with current technology that allows for detection around these thresholds. On the extensive margin, qualitative evidence suggests that public water systems are likely responding to the policy by increasing the number of treatment plants built and retrofitted to address PFAS. Some of these new plants are funded through grants and are located pri-

marily in counties with densely populated urban areas.

In the absence of an enforceable federal standard, state regulation for emerging contaminants of concern can guide federal policy. California's notification levels led the charge for policy surrounding PFAS in drinking water but, ultimately, I cannot conclude that it did or did not change concentration. However, increases in treatment capabilities signal a commitment to address contamination from these harmful chemicals. These new facilities are located in impactful areas in terms of contamination. Furthermore, California's allocation of funds to address PFAS underscores the role policymakers play in incentivizing reductions of concentration in highly polluted areas that may impact disadvantaged communities. Further work is required to quantify the impacts of PFAS drinking water policy across the United States, including behavioral responses and adaptation.

Overall, my findings point to two main takeaways for policymakers and researchers. First, how we deal with interval-censored data for policy evaluation can lead to contradictory conclusions. Deliberate analysis that accounts for this issue is critical. Second, setting thresholds that align with available technology and allow for accurate detection is essential to ensure compliance with standards set.

References

- Abrevaya, Jason and Chris Muris (2020). “Interval censored regression with fixed effects”. In: *Journal of Applied Econometrics* 35.2, pp. 198–216.
- Andarge, Tihitina et al. (2025). “Quantifying Threshold Manipulation in the Presence of Rounding: The Case of Lead Monitoring in US Drinking Water”. In: *American Economic Review: Insights* 7.3, pp. 285–305.
- Athey, Susan and Guido W Imbens (2006). “Identification and Inference in Nonlinear Difference-in-Differences Models”. In: *Econometrica* 74.2, pp. 431–497.
- Baker, Jonathan, Lori Benneer, and Sheila Olmstead (2023). “Does information disclosure reduce drinking water violations in the united states?” In: *Journal of the Association of Environmental and Resource Economists* 10.3, pp. 787–818.
- Benneer, Lori S, Katrina K Jessoe, and Sheila M Olmstead (2009). “Sampling Out: Regulatory Avoidance and the Total Coliform Rule”. In: *Environmental Science & Technology*.
- Benneer, Lori S and Sheila M Olmstead (2008). “The impacts of the “right to know”: Information disclosure and the violation of drinking water standards”. In: *Journal of Environmental Economics and Management* 56.2, pp. 117–130.
- California State Water Resources Board (2025). *Public Drinking Water Watch*. Accessed: 2025-02-26 <https://sdwis.waterboards.ca.gov/PDWW/>.
- California State Water Resources Control Board (2018-2022). *Drinking Water Quality Data*. Accessed: 2023-07-07 https://www.waterboards.ca.gov/drinking_water/certlic/drinkingwater/EDTlibrary.html.
- Chay, Kenneth Y and James L Powell (2001). “Semiparametric censored regression models”. In: *Journal of Economic Perspectives* 15.4, pp. 29–42.
- Christensen, Peter, David A. Keiser, and Gabriel E. Lade (2023). “Economic Effects of Environmental Crises: Evidence from Flint, Michigan”. In: *American Economic Journal: Economic Policy* 15.1, pp. 196–232. DOI: 10.1257/pol.20190391. URL: <https://www.aeaweb.org/articles?id=10.1257/pol.20190391>.
- Cook, Jason and James McDonald (2013). “Partially adaptive estimation of interval censored regression models”. In: *Computational Economics* 42, pp. 119–131.
- Cutler, David and Grant Miller (2005). “The role of public health improvements in health advances: the twentieth-century United States”. In: *Demography* 42.1, pp. 1–22.
- DeLuca, Nicole M et al. (2022). “Human exposure pathways to poly-and perfluoroalkyl substances (PFAS) from indoor media: A systematic review”. In: *Environment international* 162, p. 107149.

- Domingo, José L and Martí Nadal (2019). “Human exposure to per-and polyfluoroalkyl substances (PFAS) through drinking water: A review of the recent scientific literature”. In: *Environmental research* 177, p. 108648.
- Environmental Protection Agency (2018-2022). *Safe Drinking Water Information System (SDWIS) Federal Reporting Services*. Accessed: 2024-10-22 https://sdwis.epa.gov/orders/sfdw_pub/r/sfdw/sdwis_fed_reports_public/1?clear=1.
- Foster, Stephanie A et al. (2019). “Arsenic drinking water violations decreased across the United States following revision of the maximum contaminant level”. In: *Environmental science & technology* 53.19, pp. 11478–11485.
- Ghanem, Dalia, Désiré Kédagni, and Ismael Mourifié (2023). “Evaluating the impact of regulatory policies on social welfare in difference-in-difference settings”. In: *arXiv preprint arXiv:2306.04494*.
- Ghanem, Dalia, Shu Shen, and Junjie Zhang (2020). “A Censored Maximum Likelihood Approach to Quantifying Manipulation in China’s Air Pollution Data”. In: *Journal of the Association of Environmental and Resource Economists* 7.5, pp. 965–1003.
- Hadachek, Jeffrey (2025). *The cost of nitrate pollution in drinking water*. Working Paper.
- Horowitz, Joel L and Charles F Manski (1998). “Censoring of outcomes and regressors due to survey nonresponse: Identification and estimation using weights and imputations”. In: *Journal of Econometrics* 84.1, pp. 37–58.
- Jacqz, Irene, Tugba Somuncu, and John Voorheis (2024). *Fighting Fire with Fire(fighting Foam): The Long Run Effects of PFAS Use at U.S. Military Installations*. Working Paper.
- Jenkins, Stephen P et al. (2011). “Measuring inequality using censored data: a multiple-imputation approach to estimation and inference”. In: *Journal of the Royal Statistical Society Series A: Statistics in Society* 174.1, pp. 63–81.
- Keiser, David (2019). “The Missing Benefits of Clean Water and the Role of Mismeasured Pollution”. In: *Journal of the Association of Environmental and Resource Economists* 6.4, pp. 669–707.
- Keiser, David et al. (2023). “Water works: Causes and consequences of safe drinking water in America”. In: *Quarterly Journal of Economics*. Manuscript in revision.
- Keiser, David A, Sheila M Olmstead, et al. (2021). “A water rule that turns a blind eye to transboundary pollution”. In: *Science* 372.6539, pp. 241–243.
- Keiser, David A and Joseph S Shapiro (2019). “Consequences of the Clean Water Act and the demand for water quality”. In: *The Quarterly Journal of Economics* 134.1, pp. 349–396.

- Liddie, Jahred M, Laurel A Schaidler, and Elsie M Sunderland (2023). “Sociodemographic factors are associated with the abundance of PFAS sources and detection in US community water systems”. In: *Environmental science & technology* 57.21, pp. 7902–7912.
- Manski, Charles F and Elie Tamer (2002). “Inference on regressions with interval data on a regressor or outcome”. In: *Econometrica* 70.2, pp. 519–546.
- Marcus, Michelle (2022). “Testing the water: Drinking water quality, public notification, and child outcomes”. In: *Review of Economics and Statistics* 104.6, pp. 1289–1303.
- (2025). “Testing above the limit: Drinking water contamination and test scores”. In: *Journal of Public Economics* 243, p. 105313.
- Maruzzo, Aaron J et al. (2025). “Socioeconomic Disparities in Exposures to PFAS and Other Unregulated Industrial Drinking Water Contaminants in US Public Water Systems”. In: *Environmental Health Perspectives* 133.1, p. 017002.
- Meyer, Andrew and Zach Raff (2024). “The effect of water pollution regulation on prices: Evidence from Wisconsin’s phosphorus rule and sewer utility bills”. In: *Journal of Environmental Economics and Management* 124, p. 102914.
- Mueller, Rosie et al. (2024). “Quantifying disparities in per-and polyfluoroalkyl substances (PFAS) levels in drinking water from overburdened communities in New Jersey, 2019–2021”. In: *Environmental Health Perspectives* 132.4, p. 047011.
- National Academies of Sciences, Engineering and Medicine (2022). “Potential Health Effects of PFAS”. In: *Guidance on PFAS Exposure, Testing, and Clinical Follow-Up*. National Academies Press (US).
- Pennino, Michael J, Jana E Compton, and Scott G Leibowitz (2017). “Trends in drinking water nitrate violations across the United States”. In: *Environmental science & technology* 51.22, pp. 13450–13460.
- Rios-Avila, Fernando, Gustavo Canavire-Bacarreza, and Flavia Sacco-Capurro (2024). “Recovering income distribution in the presence of interval-censored data”. In: *The Journal of Economic Inequality*, pp. 1–22.
- Sunderland, Elsie M et al. (2019). “A review of the pathways of human exposure to poly- and perfluoroalkyl substances (PFASs) and present understanding of health effects”. In: *Journal of exposure science & environmental epidemiology* 29.2, pp. 131–147.
- U.S. Census Bureau (2020). “*tl_rd22_us_uac20.shp*” *TIGER/Line Shapefiles*. Accessed: 2023-02-14 <https://www.census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.2020.html>.
- (2018-2022). *American Community Survey*. Accessed: 2023-02-09 <https://data.census.gov>.

- Vo, Hoang Nhat Phong et al. (2020). “Poly-and perfluoroalkyl substances in water and wastewater: A comprehensive review from sources to remediation”. In: *Journal of Water Process Engineering* 36, p. 101393.
- Waterfield, Gina et al. (2020). “Reducing exposure to high levels of perfluorinated compounds in drinking water improves reproductive outcomes: evidence from an intervention in Minnesota”. In: *Environmental Health* 19, pp. 1–11.

Appendix

Assessing the Impact of PFAS Water Regulation

Laura Alcocer Quñones

Contents

A	Analysis for PFOA	A2
B	Robustness	A6
B.1	Sample Size	A6
B.2	Aggregation Over Multiple Sampling Points	A8
B.3	Partial Identification of Quantile Treatment Effects on the Treated	A9
C	Supplementary Figures and Tables	A14
D	Data Cleaning Description	A19
D.1	Intensive margin	A19
D.2	Extensive margin	A20

A Analysis for PFOA

For ease of exposition, the corresponding analysis for PFOA is contained in this section. As shown previously, using interval-censored data dropping censored observations or replacing them with the bounds of the interval can lead to different conclusions. As before, assuming that parallel trends holds (see Figure A10), I apply DID to quantify the impact of this policy on PFOA concentration. Table A1 shows results for naïve DID and CIC models using the approaches described. The coefficients of interest indicate the policy had no impact on water quality using these approaches.


Table A1: Average Treatment Effects for PFOA Using DID and CIC

	(1)	(2)	(3)	(4)	(5)
	Y_{it}^z	Y_{it}^r	$Y_{it}^* S_{it} = 0$	$Y_{it}^{\frac{r}{2}}$	$Y_{it}^{\sqrt{\frac{r}{2}}}$
ATT DID	-0.653 (0.938)	-2.884* (1.183)	-2.234 (1.860)	-1.808 (1.016)	-0.984 (0.938)
ATT CIC	0.036 (0.910)	-1.525 (1.060)	-0.701 (1.888)	-0.828 (0.950)	-0.152 (0.873)
Observations	10055	10055	4356	10055	10055
N-PWS	1224	1224	615	1224	1224

Notes: The unit of observations is facility-quarter-year. The sample includes facilities within PWS testing for PFAS in California and New Jersey between 2018 and the first quarter of 2020. The independent variable is an indicator for the post period. The dependent variables are the upper bound of PFOA concentration (column 1), the lower bound of PFOA concentration (column 2), PFOA concentration dropping censored observations (column 3), censored PFOA concentration replaced with half of the reporting limit (column 4) and censored PFOA concentration replaced with the square root of half the reporting limit (column 5). CIC estimates bootstrapped with 1,000 replications. Standard errors clustered at the PWS level are reported in parentheses: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Threshold policies are likely to affect different parts of the distribution in a way that is not captured by the average effect. To quantify these impacts, I use CIC to estimate the counterfactual distribution for PFOA concentration. The counterfactual and observed distributions of PFOA concentration are shown in Figure A12. Fixing a quantile, the vertical difference between these distributions yields the quantile treatment effects on the treated (QTT). Figure A1 shows the QTT across the distribution. Standard errors were bootstrapped and clustered at the public water system level. Three of the cases shows decreases in concentration for different parts of the distribution while two of them show no significant impact throughout. How interval-censored data is employed in policy evaluation can lead to different policy conclusions. This finding is consistent across contaminants.

Figure A1: Quantile Treatment Effects with Changes-in-Changes for PFOA




Notes: Each panel shows the quantile treatment effects on the treated for PFOA concentration from the 10th to the 90th. Top panel on the left corresponds to the case where censored observations are replaced with zeros (upper bound), the top middle corresponds to the case where censored observations are replaced with the reporting limit (lower bound), the top right corresponds to the case where censored observations are dropped, the bottom left corresponds to the case where the censored observations are replaced with half of the reporting limit and the bottom middle corresponds to the case when censored observations are replaced with the square root of half the reporting limit. The last two cases are common suggestions among practitioners.


To address interval censoring in drinking water samples, I employ the parametric approach described in the main body of the paper (Section 3.2). The parametric fit to a Weibull distribution for PFOA concentration in New Jersey and California is shown in Figure A2a. As for PFOS, the parametric fit for the observed portion of the distribution is adequate and falls within the 95% confidence interval.

Figure A2: Fitted PFOA Concentration Distribution for New Jersey and California in Pre and Post Periods

(a) New Jersey




(b) California



Notes: Panel (a) corresponds to New Jersey where the left figure shows the parametric fit for the time period before notification levels changed and the right panel corresponds to the fit for the time period after. Analogously, panel (b) corresponds to California. Both groups and time periods are fitted using a Weibull distribution. The lower bound (dashed line) denotes the distribution of PFOA concentration replacing censored observations with the reporting limit. The upper bound (solid line) denotes the distribution of PFOA concentration replacing censored observations with zero. The blue line corresponds to the Weibull fit for the uncensored distribution of PFOA concentration. 95% confidence intervals are provided for Weibull fit and the uncensored portion of the distribution.

PFOA concentration is highly censored in the pre-period for California, with a reporting limits up to 20 ppt. To overcome this issue, I employ the same approach as before and estimate the counterfactual distribution using the CDFs estimated using MLE. Then, I estimate quantile treatment effects on the treated for PFOA. Figure A3 the results along the distribution. I find no evidence that the policy had an impact on water quality measured by PFOA concentration. A notification level change is likely to affect the lower portion of the distribution more and the presence of reporting limits and censored observations limits our ability to detect those impacts. Furthermore, impacts are estimated using a small post-period sample that captures very short term response. Since the majority of new and retrofitted treatment plants came online after mid 2020, it is not surprising that I cannot reject the hypothesis that this policy had an impact on water quality.

Figure A3: Quantile Treatment Effects on the Treated for PFOA Concentration




B Robustness

B.1 Sample Size

One concern for the analysis presented is that the sample may be too small to perform appropriate inference. As a robustness check, I include PFAS concentration data for New York (NY) and Massachusetts (MA) in addition to New Jersey. These states are comparable to California since they also had drinking water guidance that was purely informational during the sample period. Data on these states is sparse for the years studied and reduced even further when accounting for facilities tested before and after. Overall, the new states add 1,080 observations to my sample bringing the total to 11,152 facility-quarters.

Figure A4 shows the fit to a Weibull distribution using censored maximum likelihood. The fit is significantly worse than just including New Jersey. What is interesting in this figure is the magnitude of the maximum reporting limit since it is much larger in the post period than in the pre period. The large quantity of observations around 10 ppt correspond to NY data that had faulty reporting limits in the data. To circumvent any issues with reporting, I assign the maximum reporting limit in the period (pre or post) to observations that show no concentration and a reporting limit of zero. Usually, we expect reporting limits to become smaller with time and technological advances but that may not be the case here and could be driven by testing in a specific laboratory with outdated methods.


Figure A4: Weibull Fit for PFOA and PFOS including NY and MA as Controls



Notes: Panel (a) corresponds to the results for PFOA. Each subpanel has 3 corresponding lines: the upper bound depicted by the gray solid line, the lower bound depicted by the gray dashed line and the blue line depicts the fit. As before, the data are censored at the maximum reporting level in a period. Data contained inside the yellow rectangle are censored for maximum likelihood estimation. Panel (b) corresponds to the results for PFOS.

Setting aside this issue, using these distributions in changes-in-changes leads to very similar results as found previously. The bounds of the quantile treatment effects on the treated (Figure A5) show the same pattern for both contaminants as before.

Figure A5: Results for PFOA and PFOS including NY and MA




Notes: Left panel shows the quantile treatment effects on the treated for PFOA concentration and the right panel shows the quantile treatment effects on the treated for PFOS concentration. Shaded area denotes 95% point-wise confidence intervals. Standard errors are computed through 1,000 bootstrap replications with resampling at the water system level. Control group includes New Jersey, New York and Massachusetts.

B.2 Aggregation Over Multiple Sampling Points

An additional concern is how water samples are aggregated within a treatment plant. Plants can have multiple sampling points which can be deemed problematic. The primary concern is that by taking the average over treatment plants, the readings are no longer comparable to those from raw water in wells. To assuage this concern, I take the minimum and the maximum concentration over all sampling points within a treatment plant in a given quarter. Figure A6 show the results for PFOS and PFOA respectively. These results show that how we define concentration within a treatment plant does not affect the results.

Figure A6: Results Using Maximum and Minimum Concentration Within Treatment Plants



Notes: Panels (a) through (d) show quantile treatment effects on the treated using the maximum or minimum concentration at treatment plants within a quarter for each contaminant. Shaded area denotes 95% point-wise confidence intervals. Standard errors computed through 1,000 bootstrap replications with resampling at the water system level.

B.3 Partial Identification of Quantile Treatment Effects on the Treated

The main body of the paper shows the case of point-identified estimates. However, one can partially identify QTT if the parametric fit for the treated group is inadequate or if one faces difficulty in MLE convergence for this group. For such cases, one can use the bounds of the distribution for the treated group to bound the QTT. In this context, I use the upper and lower bounds of concentration in the treatment group (California) in the pre-period as defined in equation 4. This allows me to estimate the lower bound of the counterfactual distribution through the equation:


$$H_{11}^{CF, LB} = H_{10}^r(H_{00}^{-1}(H_{01}(y))) \quad (A1)$$

and I can estimate the upper bound of the counterfactual distribution through the equation:

$$H_{11}^{CF,UB} = H_{10}^z(H_{00}^{-1}(H_{01}(y))) \quad (A2)$$

To illustrate these equations, Figure A7 shows the bounds the bounds for the observed and counterfactual PFOS distribution in California during the post period. The counterfactual estimated is PFOS concentration in California in the absence of the notification level change. The panel on the left corresponds to H_{11}^{LB} and H_{11}^{UB} . The panel on the right correspond to the estimates from equations A1 and A2. Panel A7c shows the inverse of the observed and estimated counterfactual distribution bounds.

Figure A7: Bounds for Observed and Counterfactual PFOS Distribution in California in the Post-Period



Notes: Panel (a) shows the post-period bounds for the observed distribution of PFOS concentration in California. Panel (b) corresponds to the bounds of the counterfactual distribution estimated using changes-in-changes. Panel (c) puts these two figures together and shows the inverse distribution of the bounds for the observed and the bounds for the counterfactual. This figure illustrates how quantile treatment effect bounds are estimated by taking vertical differences.

Quantile treatment effects are estimated by fixing a quantile and taking the vertical difference between the observed and the counterfactual distribution. In this case, we only have bounds for the observed and the counterfactual distributions. Define the treatment effect on the treated at quantile q as

$$\tau_q^{CIC} \equiv H_{11}^{-1}(q) - H_{11}^{-1,CF}(q)$$

where $H_{11}^{-1}(q)$ is the inverse of the distribution of observed PFAS concentration for California


in the post-period and $H_{11}^{-1,CF}(q)$ is the inverse of the counterfactual distribution of PFAS concentration for California (i.e., the inverse of the distribution in the absence of the regulatory change). Recall that we do not observe the latent distribution of concentration, but rather its bounds. Therefore, the bounds for the quantile treatment effects can be expressed as follows.

Since $H_{11}^{-1}(q) \in [H_{11}^{-1,UB}(q), H_{11}^{-1,LB}(q)]$ and $H_{11}^{-1,CF}(q) \in [H_{11}^{-1,CF,UB}(q), H_{11}^{-1,CF,LB}(q)]$, then

$$\tau_q^{CIC} \in [H_{11}^{-1,UB}(q) - H_{11}^{-1,CF,LB}(q), H_{11}^{-1,LB}(q) - H_{11}^{-1,CF,UB}(q)]$$

Figure A8 shows the bounds for the quantile treatment effects on the treated for PFOS and PFOA, respectively. The blue line indicates the upper bound, the red line indicates the lower bound and the gray area denotes the 95% point-wise confidence band on the set. Using this approach, I find that I am unable to reject the hypothesis that the policy impacted water quality. However, the purpose of this section is to provide practitioners with flexibility when dealing with interval-censored data for policy evaluation.

Figure A8: Bounds on the Quantile Treatment Effects on the Treated



Notes: The blue line corresponds to the upper bound and the red line corresponds to the lower bound of the quantile treatment effects on the treated. Gray area denotes the 95% point-wise confidence intervals on the QTT bounds. Standard errors are computed using 1,000 bootstrap replications with resampling conducted at the water system level.


C Supplementary Figures and Tables

Table A2: Number of Facilities in California By Status and Monitoring Schedule Change Eligibility

	PFOA		PFOS	
	Eligible for Schedule Change	Not Eligible for Schedule Change	Eligible for Schedule Change	Not Eligible for Schedule Change
Active Facility	520	395	471	442
Inactive Facility	33	33	32	36


Notes: Table shows count of facilities in California that tested before and after notification level change from 2018 through the first quarter of 2020. Facilities are classified by activity status and eligibility for monitoring schedule changes. Eligibility was determined using the previous 4 quarters of results for each pollutant and activity status was provided in the sampling data for each facility.

Figure A9: Parallel Trends for PFOS




Notes: Average PFOS concentration per period by group. Each panel corresponds to a different definition of concentration. Time period spans from 2018 through the first quarter of 2021 to avoid confounding effects with response level changes.

Figure A10: Parallel Trends for PFOA




Notes: Average PFOA concentration per period by group. Each panel corresponds to a different definition of concentration. Time period spans from 2018 through the first quarter of 2021 to avoid confounding effects with response level changes.

Figure A11: Observed vs. Counterfactual Quantile Distribution of PFOS in California



Notes: Each panel can be interpreted by fixing a quantile and taking the vertical difference between the observed concentration and the counterfactual concentration, meaning the vertical difference between the blue and the red line. Standard errors for the counterfactual distribution are bootstrapped with resampling at the PWS level.

Figure A12: Observed vs. Counterfactual Quantile Distribution of PFOA in California




Notes: Each panel can be interpreted by fixing a quantile and taking the vertical difference between the observed concentration and the counterfactual concentration, meaning the vertical difference between the blue and the red line. Standard errors for the counterfactual distribution are bootstrapped with resampling at the PWS level.

Table A3: Summary Statistics for Extensive Margin: 2018-2022

	Mean	SD	Min	Max	Observations
Water Quality					
Cens. PFOA Replaced w/Zero (ppt)	5.38	16.27	0	350	14,398
Cens. PFOS Replaced w/Zero (ppt)	9.88	25.20	0	510	14,399
Cens. PFOA Replaced w/RL (ppt)	6.772	15.964	0.002	350	14,398
Cens. PFOS Replaced w/RL (ppt)	11.348	24.959	0.002	510	14,399
PFOA Reporting Limit (ppt)	2.442	2.420	0.002	20	8,040
PFOS Reporting Limit (ppt)	2.950	5.338	0.002	40	7,034
Population Served by PWS (thousands)	173.10	327.77	0	3,806.27	14,400
County Characteristics					
Num. of Households (thousands)	1,284.08	1,262.21	23.95	3,415.73	14,251
Average HH Income (USD)	112,123.51	27,076.21	64,050.00	216,720	14,251
Total Population (thousands)	3,813.6	3,674.8	64.4	10,105.5	14,251
Num. of Housing Units (thousands)	986.9	1,276.5	0.9	3,620.2	14,251
% White Population	47.3	15.2	29.4	93.1	14,251
% Black Population	5.7	2.8	0.2	13.6	14,251
% Hispanic Population	40.7	12.6	8.6	85.0	14,251
% Urban Area	0.9	0.4	0	1.0	14,400

Notes: The unit of observation is a facility-quarter. PFAS concentration data are from the California State Water Board, socioeconomic and demographic characteristics are from the U.S. Census Bureau. The sample is restricted to wells and treatment plants in California that tested for PFAS between 2018 and 2020. Reporting limits are provided for readings with censored concentration. Multiple water systems are assigned the same county characteristics based on principal county served.

Figure A13: PFAS Concentration and Population Served for Treatment Plants After Policy Change



Notes: Population served is sourced from SDWIS and matched with concentration data based on PWSID. Each point represents a facility-quarter observation for treatment plants after the policy change. Blue dots correspond to existing plants, red diamonds correspond to new treatment plants and yellow diamonds correspond to retrofits. Concentration was winsorized at 100 ppt.

Table A4: Socioeconomics and Demographic Characteristics for Counties with Newly Active vs. Existing Wells

	(1) Newly Active Wells	(2) Existing Wells	(3) Difference (1)-(2)
Num. of Households (thousands)	1465.1 (1286.6)	1263.8 (1261.6)	201.34** (0.013)
Average HH Income (USD)	117517.6 (31106.8)	112864.5 (27854.4)	4653.10** (0.017)
Total Population (thousands)	4380.0 (3703.1)	3751.1 (3680.1)	628.92*** (0.007)
Population Served by PWS (thousands)	236.7 (308.8)	176.1 (346.2)	60.61*** (0.002)
Num. of Housing Units (thousands)	1033.1 (1315.7)	987.2 (1280.4)	45.92 (0.577)
% White Population	43.9 (14.3)	47.6 (15.3)	-3.65*** (0.000)
% Black Population	6.1 (2.4)	5.6 (2.9)	0.46*** (0.003)
% Hispanic Population	41.9 (10.9)	40.2 (12.9)	1.67** (0.015)
% Urban Area	0.8 (0.4)	0.8 (0.4)	-0.04* (0.081)
Observations	261	12,266	12,527

Notes: Table shows difference in means for demographic and socioeconomic characteristics between counties with newly active or existing wells. Column (1) corresponds to counties with PWS that had newly active wells and column (2) corresponds to counties with PWS that had existing wells. Data corresponds to annual ACS matched with the extensive margin panel using principal county served by a public water system. Standard deviations are reported for the means and p-values are reported for the differences: * $p < 0.1$; ** $p < 0.05$; ***

D Data Cleaning Description

D.1 Intensive margin

For the intensive margin analysis, I use raw data on PFAS concentration in drinking water sourced from the California State Water Resources Control Board and the New Jersey Department of Environmental Protection. The former is publicly available for download and the latter was obtained through a public records request. Each observation corresponds to a test in a sampling point for a PFAS contaminant. Facilities may have more than one sampling point, as is often the case for treatment plants. I subset the panel for tests corresponding to PFOA and PFOS from 2018 through the first quarter of 2020. I keep tests for wells and treatment plants, as these types of facilities are where most of PFAS testing occurs in California. For each test, I create two new variables: (1) concentration with censored observations replaced with zeros and (2) concentration with censored observations replaced with the reporting limit. Some of these tests have an error in reporting where there is no concentration provided, and the reporting limit is set to zero. Reporting limits are always positive. To address this issue, I input the maximum reporting limit over a period (before or after the policy) by state for those cases. This allows me to provide a conservative estimate

for that test.

These measurements are then aggregated to the quarter level by facility, resulting in average PFAS concentration per facility-quarter. Because public water systems can implement treatment in response to the policy, taking the average concentration over raw and finished water allows me to uncover the effects of this policy prior to distribution to consumers. Additionally, I take the maximum and minimum concentrations at treatment plants.

D.2 Extensive margin

For the extensive margin analysis, I assemble a comprehensive panel of facilities for public water systems testing for PFAS in California by quarter from the Safe Drinking Water Information System (SDWIS).¹ I merge this panel using state facility ids with the testing data to find when facilities, in my case wells and treatment plants, go inactive, disappear from SDWIS, are added for the first time, and the timing and type of treatment implemented.

To find when wells or treatment plants drop out of the SDWIS database, I create an indicator variable when a facility has been observed for the last time using a backbone with all facilities reporting every period. This backbone provides a comparison to find if facilities stop reporting. A facility is considered as dropped when it does not show up in SDWIS for 4 consecutive periods. To check for cases where a well is taken offline for extended periods of time, I create a variable to identify the last time a facility is tested and fix any observations that are double counted as dropped. Then, I sum of the number of facilities that drop out for each type of facility in each quarter. Additionally, I find the number of facilities that become inactive and are untested after notification levels changed. This case differs from the drop out facilities because I have the exact date of inactivity and facilities do not show up for the remainder of the period studied. In my whole sample, there is a single case when a facility is made inactive and does not test or appear in sample after the policy takes place.

To find when new wells or treatment plants are added, I create an indicator variable when a facility is observed for the first time using a backbone with all facilities reporting every period. This backbone provides a comparison to find if facilities are added to SDWIS. I create an indicator for when a facility has been tested and was missing for 3 consecutive periods prior to that test. To avoid double counting facilities, I identify the first time each of these were tested and eliminate duplicates if they show multiple time periods. In addition

¹This information can be obtained from detailed reports on facilities specifying public water system identifiers and dates.

to this method, I spot checked cases where the count of new wells appeared to be high and compared them to state records. For treatment plants, I classify them as new or retrofitted based on state and federal data. If a treatment plant appears in SDWIS the prior period with a specific name and the following period has that name but the string “PFAS” included, then that plant is considered a retrofit. Otherwise, treatment plants are assigned as new facilities. For wells, I check SDWIS data from 2016 through the end of 2022 to make sure these wells are not appearing before then. Unfortunately, this approach has limitations since wells can go untested for long periods of time.

As an additional check, I use well completion reports in the state of California for new wells for public water supply. This data provides geographic information for new wells constructed across the state. I cannot match specific well identifiers with sampling point identifiers due to data constraints. Instead, I perform a raw count of newly finished wells for public water systems that test for PFAS over time. Since the database considers all wells, it is expected that my sample is a subset of the total reported by the state. The count of new wells using this method is larger than my own, providing reassurance that the count of “new” wells identified from SDWIS is reasonable.

To explore the relationship between these new facilities and demographic/socioeconomic characteristics, I merge geospatial data on public water system boundaries obtained from the Department of Drinking Water with data on new wells and treatment plants, as well as yearly data at the county level from the American Community Survey (ACS) for 2018 through 2022 provided by the Census Bureau. ACS data was merged on primary county served provided by SDWIS. Additionally, I intersect PWS boundaries to the urban-rural shapefile provided by the Census Bureau. This allows me to assign PWS as urban or rural based on a binary variable.