Know Thy Pipes: How Disclosure of Proxy Information Shapes Household Demand for Waterborne Lead Mitigation and Exposure Assessments*

Fanyu Wang† Siyuan Hu‡

March 23, 2026

Preliminary draft. Please do not cite without permission.
Most updated version here.

Abstract

We study how individualized, proxy-based environmental information shapes beliefs, defensive investments, and demand for precise diagnostics in the context of waterborne lead exposure. Leveraging the mandated disclosure of water service line inventories under the 2021 Lead and Copper Rule Revision, we conduct a pre-registered field experiment in four large U.S. cities that discloses to treated households whether their building's service line is recorded as lead. Disclosure induces sharp belief updating. Among homes recorded as having lead service lines (LSLs), perceived LSL probability rises by 21 percentage points relative to non-lead homes; meanwhile, belief confidence increases by 0.2 points (on a 1-5 scale) across all treated respondents. Willingness to pay for a pitcher-style lead filter falls by $8 on average among the non-lead respondents, while the effect on the lead respondents is ambiguous. At the same time, willingness to pay for a professional tap-water lead inspection increases by $9 among lead respondents and decreases by $7 among non-lead respondents. We develop a stylized model in which households are uncertain about latent pipe material and about realized tap-water contamination conditional on that material. Disclosure resolves the former but not the latter, so mitigation loads on the between-state contrast in expected damages, whereas diagnostics load on the residual within-state uncertainty. This framework rationalizes the asymmetric responses we document. Back-of-the-envelope calculations suggest that disclosure of water service line material type alone could generate approximately $17 million in social value across our four study cities, even before accounting for any health gains from subsequent mitigation practices.

JEL Codes: C93, D83, I18, Q53, Q58

*We would like to thank Douglas Almond, Joel Ferguson, Corbett Grainger, Rhiannon Jerch, Gabriel Lade, Robert Metcalfe, Daniel Phaneuf, Jeffrey Shrader, Christopher Timmins, and seminar participants at Columbia University, the Eighth Annual Social Cost of Water Pollution Workshop, the J-TREE Regional Workshop, Second Women in Environmental Economics Workshop, University of Wisconsin-Madison, and the WEAI Annual Conference for their helpful comments. We gratefully acknowledge the financial support from the CEEP Student Research Grants, the CELSS Graduate Student Experimental Project Grant at Columbia University, and the AAE Department Research Awards at the University of Wisconsin-Madison. Gezi Fang, Geraldine Hu, and Navya Kumar provided excellent research assistance. All mistakes are our own.

†School of International and Public Affairs, Columbia University. Email: fw2397@columbia.edu.

‡Department of Agricultural and Applied Economics, University of Wisconsin-Madison. Email: shu248@wisc.edu.

1 Introduction

Lead exposure through drinking water remains a first-order public health concern in the United States. The principal pathway is corrosion from lead service lines (LSLs)—the pipes made of lead and connecting buildings' internal plumbing to municipal mains. Although Congress prohibited their use in public water systems in 1986, roughly nine million LSLs still serve households nationwide as of 2023 (USEPA, 2023). Lead is a potent neurotoxin with no safe level of exposure; there are no effective secondary interventions (Bui et al., 2024). Upon entry, the human body falsely recognizes lead as similarly charged (+2) metallic elements, like calcium or iron, causing disruptions in many critical processes throughout the body (Garza et al., 2006). For example, lead can cross the blood-brain barrier, interfering with neurotransmitter functions and spurring cognitive changes. For children in particular, even low doses impair cognition and behavior with persistent consequences for schooling performance, adult mental health, and socioemotional development (Aizer et al., 2018; Aizer and Currie, 2019; Grönqvist et al., 2020; Hollingsworth et al., 2022; Reyes, 2015). Lead exposure is also linked to diminished fertility (Grossman and Slusky, 2019), adverse birth outcomes (Clay et al., 2014, 2025; Dave and Yang, 2022), deficits in cognitive skills (Ferrie et al., 2012; Marcus, 2023), and overall mortality (Hollingsworth and Rudik, 2021; Menke et al., 2006). Addressing exposure at the household margin is, therefore, central to the welfare calculus of U.S. water infrastructure investments and environmental policymaking.

This paper provides the first experimental evidence on how individualized but proxy-based environmental information shapes beliefs, defensive investments, and the demand for more accurate exposure assessment, in the context of waterborne lead. We design and implement a pre-registered, within-city randomized information intervention in four large U.S. cities—Detroit, Indianapolis, Milwaukee, and New York City (NYC).[1] Respondents assigned to the treatment group learn about the LSL status of their residential buildings according to the city inventories. Before and after disclosure, we elicit (1) belief about having an LSL and the perceived number of lead exceedances in 100 kitchen tap water samples, along with belief confidence, and (2) incentive-compatible willingness to pay (WTP) for a pitcher-style lead filter and for a professional tap-water lead inspection.

Information disclosure is often viewed as a scalable, relatively low-cost tool to incentivize private

[1] This experiment is registered in the AEA RCT Registry as AEARCTR-0016286.

defensive behavior. U.S./ EPA mandates individualized disclosure of LSL status under Lead and Copper Rule Revision (LCRR), which requires utilities to create and publicly publish parcel-level service-line inventories (USEPA, 2020). Yet two core uncertainties limit our ability to predict its effectiveness in this setting. First, an inventory disclosure is proxy-based: it informs households about pipe material rather than directly measuring tap-water lead, which depends on corrosion control, pipe condition, premise plumbing, and usage patterns. Second, even when information is precise about the proxy, its behavioral impact depends on how households learn—how disclosure shifts the perception of exposure (i.e., its mean and standard deviation) and how those beliefs map into demand for defensive investments in stopgap mitigation or precise diagnostics to resolve the remaining uncertainty.

Our design is guided by a stylized framework in which households are uncertain about both a latent infrastructure state and realized tap-water contamination. Let $S \in \{0, 1\}$ denote true service-line material, where $S = 1$ indicates a lead line, and let W denote realized contamination at the tap. Households begin with a subjective prior $q = \Pr(S = 1)$, which may exceed the objective prevalence when lead risk is salient. Conditional on $S = s$, contamination is drawn from a distribution F_s. Disclosure reveals S, but not W. This distinction matters for the two goods we study. For a filter, WTP depends on the expected damage reduction under the posterior over W, so the response to disclosure is governed by the difference between the conditional contamination distributions F_1 and F_0. If households begin overly pessimistic and line material only loosely maps into realized contamination, learning $S = 0$ can sharply reduce filter demand, whereas learning $S = 1$ need not raise it much. For an inspection, WTP equals the ex-ante value of learning W after the proxy state has already been revealed. Inspection demand is therefore positive only when residual within-state uncertainty remains large enough that direct measurement could change subsequent mitigation choices. Thus, our model highlights that proxy-based disclosure can reallocate demand between mitigation and diagnostics, even when it does not lead to a strong increase in all defensive spending. We take these predictions to the data using intent-to-treat (ITT) specifications with city fixed effects, complemented by (1) a "post-on-pre" formulation and (2) variants that replace true LSL status with a perception gap regressor to study updating rates.

As fielding of the survey is still ongoing, we expect to reach a final sample of approximately 1,600 respondents by March 2026. With the current sample of 568 valid responses from Detroit,

Indianapolis, Milwaukee, and NYC, two main findings have begun to emerge. First, disclosure induces sharp belief updating. Among respondents in homes recorded as lead, the perceived probability of having an LSL rises by about 21 percentage points relative to non-lead homes. Across all treated respondents, belief confidence increases by roughly 0.22 on a 1–5 scale, suggesting that disclosure narrows subjective uncertainty around the tap-water lead content level. Second, the results indicate a differential response across generic mitigation and precise diagnostics. Filter WTP falls by \$8 on average for the non-lead respondents ($\sim$ 13% of baseline), while the effect for the lead respondents is statistically imprecise and centered near zero. Meanwhile, inspection WTP shows an average increase of \$9 for the lead respondents ($\sim$ 13% of baseline) and an average decrease of \$7 for the non-lead respondents ($\sim$ 10% of baseline). We interpret this pattern as consistent with households treating service-line material as a noisy upstream risk signal rather than as a sufficient statistic for realized tap-water contamination. For non-lead households, disclosure corrects overly pessimistic priors and lowers demand for both stopgap mitigation and further diagnostics. For lead households, disclosure conveys that the infrastructure is riskier, but it does not, by itself, resolve how much lead is present at the tap. As a result, demand shifts more strongly toward precise diagnostics than toward a stopgap mitigation good such as a pitcher filter. These magnitudes are economically meaningful and are robust across alternative specifications, including post-on-pre outcomes and models that specifically estimate learning rates using the prior-perception gap.

The policy implications are direct—we inform the optimal design of the costliest U.S. water infrastructure intervention in decades. Lead exposure via drinking water infrastructure has remained an important piece of hidden information for the entire history of modern U.S./, contributing to disparities in pollution exposure across socio-economic groups (Hausman and Stolper, 2021). Our findings show that households substantially overestimate the probability that their own service line is made of lead—the true LSL share in our sample is about 29%, whereas average priors are around 57%. The large chasm between beliefs and reality stands in contrast to many settings in which households underestimate the severity of ambient pollutants, such as $PM_{2.5}$ (Ahmad et al., 2022; Chowdhury et al., 2025; Greenstone et al., 2021). This asymmetry makes water infrastructure a particularly informative setting for studying how individualized, proxy-based information re-aligns beliefs and defensive investments. With Congress's allocation of \$15 billion for LSL replacement in the 2021 Infrastructure Investment and Jobs Act and the EPA's mandate to disclose invento-

ries, understanding household behavioral responses is crucial for designing cost-effective bundles of policies as well as the welfare analysis of said policies. We show that disclosure both generates efficiency gains and reveals latent demand for diagnostics among lead households, netting a social value of $\sim$ \$17 million in our sample cities alone.

Our paper contributes to several strands of literature. First, we extend the environmental information disclosure literature by causally identifying how public, proxy-based information shifts demand for more precise exposure assessment. Prior work establishes that environmental information disclosure drives avoidance behaviors and reduces exposure. Air quality warnings (Barwick et al., 2024; Gao et al., 2023; Neidell, 2010; Zivin and Neidell, 2009), water contamination alerts (Dupas et al., 2023; Madajewicz et al., 2007), environmental risk disclosures (Bakkensen and Ma, 2020; Ma et al., 2024), and health risk disclosures (Dupas, 2011) all induce welfare-improving defensive behaviors and investment. Although in many cases the publicly available information provides only a coarse proxy for an individual's true environmental exposure, far less is known about whether and how such proxy-based disclosures motivate households to seek more precise exposure assessments. Our study fills this gap by directly showing that the LSL signal triggers higher WTP for diagnostics.

Second, we build on the literature examining personalized information-provision. Greenstone et al. (2021) and Metcalfe and Roth (2025) study how individualized indoor air quality readings influence defensive behaviors and investment; Conell-Price and Mulder (2024) study belief updating pertinent to property-level flood risks; and Allcott (2011), Allcott and Rogers (2014), and Jessoe and Rapson (2014) study how household-level energy-use feedback influences residential energy consumption. Relative to these costly, individualized interventions, we evaluate a policy already deployed at scale (mandated city inventories) and show that it induces households to self-select into precise diagnostics when the disclosed signal implies a high level of exposure. To implement the personalized information provision, we conduct geospatial screening and matching, the details of which will be explained in Section 2.

Third, we contribute revealed-preference evidence from a high-stakes environmental health setting to the long-standing literature on the value of information (Stigler, 1961). Our approach is closely related to randomized information experiments that study how people update expectations when provided with objective signals under Bayesian learning (Armantier et al., 2016; Möbius

et al., 2022). A growing theoretical literature also emphasizes that information policy design requires understanding both how signals update beliefs and how changes in belief map to behavioral margins (Bergemann and Morris, 2019; Caplin and Dean, 2015). We operationalize this by jointly measuring belief distributions, confidence, and incentive-compatible valuations for both mitigation and diagnostics. Our findings overall illustrate an economically meaningful portfolio reallocation that purely belief-based designs would miss.

Finally, our results connect to the extensive literature on health and human capital consequences of lead exposure (e.g., Clay et al., 2014; Dave and Yang, 2022; Hollingsworth et al., 2022; Marcus, 2023), providing the behavioral counterpart needed to assess the benefits of disclosure and the welfare returns to complementary investments such as inspection and eventual pipe replacement. Recent work has attempted to estimate large welfare gains from the Clean Water Act and Safe Drinking Water Act (Keiser and Shapiro, 2019; Keiser et al., 2023). We provide the evidence on behavioral responses needed to evaluate policies designed to combat waterborne lead exposure.

The remainder of the paper proceeds as follows. Section 2 describes the experimental design, data, and implementation, including geospatial screening and the disclosure messages. Section 3 presents the stylized model and derives testable predictions. Section 4 introduces data and survey participants. Section 5 reports the main experimental results on belief updating, confidence, and WTP for mitigation and diagnostics, with robustness and heterogeneity. Section 6 concludes.

2 Experiment Design

We now describe the experimental design outlined in Section 1.

2.1 The LSL Information Intervention Experiment

Under EPA's LCRR, water utilities must assemble and publish parcel-level inventories of service-line materials by October 16, 2024. We leverage these newly released inventories to conduct an information provision experiment embedded in an online survey, for which Figure 1 summarizes the structure.[2]

We field our survey in four large U.S. cities—Detroit, Indianapolis, Milwaukee, and NYC—using

[2]The complete survey instrument is available here here.

Qualtrics software and implemented via Managed Research on CloudResearch, a survey provider widely used in the social sciences (Bursztyn et al., 2023; Litman and Robinson, 2020). Cities were selected for policy relevance (i.e., large LSL stocks and active inventories) and operational feasibility (Table 1). The target sample comprises $N \approx 2{,}000$ validated completions, with approximately equal allocation across treatment and control within each city. Stratified randomization at the city level ensures balanced assignment and facilitates the inclusion of city fixed effects in the analysis.

At the beginning of the survey, after screening out those who self-reported that they did not use tap water as their primary water source, participants are assigned a water service line material type through a two-step geospatial matching process. First, participants pinpoint their residence on an embedded interactive map using Qualtrics-Map with Mapbox. Then, the latitude and longitude coordinates from the pinned location are automatically cross-referenced with the geospatial water service line inventory dataset for the corresponding city, hosted on ArcGIS Online via an API. If no match is found, the participant is screened out. If a match is found, they continue with the survey.

Randomization of the information disclosure occurs only after a successful match and completion of all pre-intervention questions. Within each city stratum, eligible respondents are assigned to a control or treatment arm with equal probability. The intervention comprises a concise, standardized disclosure. The control arm received a neutral description of the city inventory database and general information about LSLs. The treatment arm received the same content, along with their building-specific LSL status, explicitly attributed to the city database. Equalizing the generic text across arms isolates the causal effect of individualized disclosure from generic salience of LSL issues. Figure 2 shows screenshots of the intervention messages. All key outcomes—beliefs about LSL presence, beliefs about tap-water samples exceeding the federal limit for lead concentration, confidence, and WTP for a filter and for an inspection—are elicited both before and after this information screen.

Beliefs include (1) the perceived probability $(0 - 100)$ that the home has an LSL and (2) the expected number of lead exceedances in 100 hypothetical tap samples. Lead exceedance is defined as lead concentration surpassing the federal limit of 0.01mg/L. Respondents also report a $1 - 5$ confidence rating for the exceedance belief. We then elicit incentive-compatible WTP for (1) a pitcher-style filter and (2) a professional kitchen tap-water inspection using the Becker–DeGroot–Marschak (BDM) mechanism (Becker et al., 1964). We implement the BDM

mechanism by randomly selecting two respondents for a binding endowment (a $150 gift card), then drawing a random posted price from $0 to $150. The winners either exchange part of the endowment ($150 − stated WTP) for the chosen good when stated WTP $\geq$ price or receive the entire $150 endowment when stated WTP $<$ price. Prior to bidding, the instrument also includes a comprehension check. Together, our implementation maintains incentive compatibility while keeping implementation costs predictable. We also randomized the order of the two WTP questions within each stratum to assess and mitigate potential order effects in stated valuations, in particular whether eliciting diagnostic demand first influences stated willingness to pay for immediate mitigation, or vice versa.

Two features of our experiment warrant additional emphasis. First, by conducting eligibility screening through geospatial matching before randomization, we ensure that all participants in both treatment and control arms are successfully linked to official government service line records. This design avoids the problem in which treated participants cannot receive building-specific information because the address lookup occurs after randomization and no match is found. This also allows us to calculate the prior perception gap for all participants. Second, the disclosure text and user interface are harmonized across cities to the extent possible; any city-specific elements are also standardized in length and tone, with city fixed effects accounting for residual heterogeneity.

We considered the possibility that our information intervention might be less effective because utilities are already notifying customers about service line materials under the LCRR. To investigate this, we asked respondents if they had been notified and if they could recall the specific information. As shown in Table A1, 29% stated they could recall a notification, and only 17% correctly recalled whether their service line was lead or non-lead. More importantly, Figure A1 shows that the status of notification and recall is not associated with the distribution of the pre-intervention perceived likelihood of LSL presence.

2.2 Empirical Specifications

The empirical specifications follow the approach in the information intervention experiment literature (Haaland et al., 2023). Our estimands are ITT effects of individualized disclosure on belief updating and demand for mitigation and diagnostics. All specifications follow the registered pre-analysis plan. Let i index respondents, and $c(i)$ denote their city. Outcomes are measured

immediately before and after the intervention message:

$$Y \in \{\text{LSL prob } (P),\ \text{Exceedances } (S),\ \text{Confidence } (C),\ \text{WTP(Filter)},\ \text{WTP(Inspection)}\}.$$

Let $L_i \in \{0,1\}$ indicate whether the city inventory records the respondent's service line as lead; $D_i \in \{0,1\}$ indicate assignment to building-specific disclosure; and $\gamma_{c(i)}$ be city fixed effects. We first establish a baseline ITT specification to model the pre-post changes:

$$\Delta Y_i \equiv \left(Y_i^{\text{post}} - Y_i^{\text{pre}}\right) = \beta_1 D_i + \beta_2 L_i + \beta_3 (D_i \times L_i) + \gamma_{c(i)} + \varepsilon_i. \tag{1}$$

Here, β_1 is the disclosure effect for non-lead homes while β_3 captures the differential disclosure effect for lead versus non-lead homes.

To improve precision and guard against regression-to-the-mean concerns, we also estimate a post-on-pre form:

$$Y_i^{\text{post}} = \theta_1 D_i + \theta_2 L_i + \theta_3 (D_i \times L_i) + \rho\, Y_i^{\text{pre}} + \gamma_{c(i)} + u_i. \tag{2}$$

Under randomization, θ_1 and θ_3 identify the same causal parameters as in (1) while ρ absorbs baseline variation and typically yields tighter confidence intervals.

To study belief updating directly, we also replace L_i with the perception gap

$$G_i \equiv L_i - P_i^{\text{pre}} \in [-1, 1],$$

where P_i^{pre} denotes the respondent's pre-treatment belief about having an LSL, rescaled from the elicited 0–100 probability to the $[0,1]$ interval. We then estimate

$$\Delta Y_i = \alpha_1 D_i + \alpha_2 G_i + \alpha_3 (D_i \times G_i) + \gamma_{c(i)} + e_i, \tag{3}$$

where α_3 is interpretable as a reduced-form "updating rate" (i.e., how strongly disclosure moves outcomes conditional on the size and direction of the prior–record discrepancy). This parametrization mirrors the expectations-formation literature, which models belief updating as a function of

prior forecast errors and recovers "updating rates" from the slope of the update on that error (e.g., Armantier et al. (2016)). In our setting, a positive and sizable α_3 is evidence that disclosure moves beliefs in proportion to the prior misperception (rather than simply shifting levels via generic salience or demand effects), and thus provides a simple test of whether updating behaves in a broadly Bayesian manner. We also consider $|G_i|$ to isolate responses to uncertainty magnitude irrespective of sign.

Section 5.4 specifies a set of heterogeneity analyses we expect to conduct when we finish fieldwork with approximately 1,600 responses.

3 Model

This section develops a compact framework that links proxy-based disclosure to households' WTP for two distinct defensive margins: (1) a stopgap mitigation technology and (2) a precise exposure assessment. The model is deliberately minimal. Its purpose is to organize the empirical interpretation in an applied setting, not to provide a standalone theory of household water demand.

The key primitive is a latent infrastructure state. Let $S \in \{0, 1\}$ denote service-line material, where $S = 1$ corresponds to a lead service line and $S = 0$ to a non-lead line. Households do not observe S ex ante and instead hold a subjective prior

$$q \equiv \Pr(S = 1).$$

Because waterborne lead crises (e.g., Flint, Michigan) are highly salient public events, we allow for misconceptions about lead risk by not requiring q to equal the objective prevalence (Bordalo et al. (2022)). Conditional on $S = s$ (information disclosure), realized tap-water contamination W is drawn from a state-contingent distribution F_s, with mean μ_s and variance τ_s^2. The disclosure reveals S but not W. Hence, before disclosure, beliefs over realized contamination are given by the mixture

$$\pi^0 = qF_1 + (1 - q)F_0,$$

whereas after disclosure of $S = s$, beliefs become

$$\pi^s = F_s.$$

This structure creates the distinction between the two goods we study. A filter is a stopgap mitigation technology: its value depends on how the disclosed infrastructure state changes expected damages through the posterior distribution of W. An inspection is an information good: its value depends on the residual uncertainty about W that remains after the state has been disclosed. Accordingly, filter demand loads on the *between-state* contrast between F_1 and F_0, whereas inspection demand loads on the *within-state* uncertainty that remains inside a given F_s.

The model therefore delivers a parsimonious interpretation of our main empirical asymmetry. If households begin with pessimistic priors about living with a lead line, non-lead disclosure can sharply reduce filter demand. By contrast, lead disclosure need not generate a large increase in filter demand if service-line material is only a noisy predictor of realized contamination. At the same time, inspection demand can rise after lead disclosure if the risky state leaves households with substantial action-relevant uncertainty about realized exposure.

The remainder of the section proceeds in three steps. We first present a general choice environment and derive the optimal BDM bid for a filter and for an inspection—objects that coincide with the experimentally elicited WTPs. We then impose quasi-linearity to obtain transparent expressions that map these WTP objects into the latent-state structure above. Finally, we specialize primitives to obtain comparative statics that we take to the data. The proofs for all propositions, lemmas, and corollaries are relegated to the Appendix.

3.1 General Setup

Let $u(x, d)$ denote household utility, where x is the numeraire with price 1 and d is the contamination of water intake. We assume baseline income y is sufficiently large (e.g., $y > \bar{r}$) to abstract from liquidity constraints and ensure interior bidding solutions.

Assumption 1. *u is concave in (x, d), twice differentiable, and satisfies $\partial u/\partial x > 0$ and $\partial u/\partial d < 0$.*

Income is denoted by y. Throughout, π denotes an arbitrary belief distribution over realized contamination W; in the application, the relevant beliefs are π^0 before disclosure and π^s after

disclosure of $S = s$. We study the implications of belief changes for the WTP for (1) a filter and (2) a professional inspection, both elicited via incentive-compatible BDM bids. For completeness, Appendix B.4 reports the benchmark in which W is known.

A filter attenuates exposure mechanically: contamination falls from W to eW with $e \in [0, 1)$ known to the household after purchase. Let bids be restricted to $[0, \bar{r}]$. If the posted BDM price $r \sim U([0, \bar{r}])$ satisfies $z \geq r$, the household buys at price r. Expected utility is therefore

$$U_F(z, \pi) = \frac{1}{\bar{r}} \int_0^z \mathbb{E}_\pi[u(y - r, eW)]\, dr + \left(1 - \frac{z}{\bar{r}}\right) \mathbb{E}_\pi[u(y, W)]\,. \tag{4}$$

Proposition 1. *Under Assumption 1, any interior optimal bid $z^*(\pi)$ for the filter satisfies*

$$\mathbb{E}_\pi\big[u\big(y - z^*(\pi), eW\big)\big] = \mathbb{E}_\pi[u(y, W)]\,. \tag{5}$$

An inspection changes information rather than exposure. Upon purchasing at price r, the household observes a realization of W before choosing mitigation effort $a \in \mathcal{A}$, which reduces contamination according to $d(W, a)$ at cost $c(a)$. Without inspection, it chooses a under the prior π. With inspection, it chooses a after observing W. Define

$$V_0(\pi) = \max_{a \in \mathcal{A}} \mathbb{E}_\pi\big[u\big(y - c(a), d(W, a)\big)\big]$$

and

$$V_1(r, \pi) = \mathbb{E}_\pi\left[\max_{a \in \mathcal{A}} u\big(y - c(a) - r, d(W, a)\big)\right].$$

The expected utility of bidding z for an inspection is therefore

$$U_I(z, \pi) = \frac{1}{\bar{r}} \int_0^z V_1(r, \pi)\, dr + \left(1 - \frac{z}{\bar{r}}\right) V_0(\pi). \tag{6}$$

Proposition 2. *Under Assumption 1, any interior optimal bid $\tilde{z}^*(\pi)$ for the inspection satisfies*

$$V_1\big(\tilde{z}^*(\pi), \pi\big) = V_0(\pi). \tag{7}$$

3.2 Comparative Statics

To obtain transparent expressions, we impose quasi-linearity:

Assumption 2. *There exists a twice-differentiable function v such that*

$$u(x, d) = x + v(d),$$

with $v'(d) < 0$.

3.2.1 WTP for Water Filter under Simplified Assumptions

Under Assumption 2, Proposition 1 implies that the optimal filter bid equals the expected reduction in contamination damages:

$$z^*(\pi) = \mathbb{E}_\pi[v(eW) - v(W)]\,. \tag{8}$$

Define

$$m(w) \equiv v(ew) - v(w), \qquad M_s \equiv \mathbb{E}[m(W) \mid S = s]\,.$$

Proposition 3. *Under Assumption 2, pre-disclosure filter WTP is*

$$z_F^0 = qM_1 + (1 - q)M_0,$$

and post-disclosure filter WTP after learning $S = s$ is

$$z_F^1(s) = M_s.$$

Hence the disclosure effects are

$$\Delta_F(1) \equiv z_F^1(1) - z_F^0 = (1 - q)(M_1 - M_0),$$

$$\Delta_F(0) \equiv z_F^1(0) - z_F^0 = -q(M_1 - M_0).$$

Proposition 3 yields the key filter intuition. If the filter is more valuable in the lead state than in the non-lead state, so that $M_1 > M_0$, then a non-lead disclosure reduces filter WTP whereas

a lead disclosure increases it. The two effects need not be symmetric. If households begin with pessimistic priors about living with a lead line, so that q is high, then the absolute decline after a non-lead disclosure can exceed the absolute increase after a lead disclosure. Moreover, if service-line material only loosely maps into realized contamination, then F_1 and F_0 are not far apart, $M_1 - M_0$ is small, and the lead-side filter response can be muted.

Corollary 1. *Under Assumption 2, if there exists $\beta > 0$ such that $v(d) = -\beta d$, then*

$$M_s = \beta(1-e)\mu_s, \qquad \mu_s \equiv \mathbb{E}[W \mid S = s].$$

Therefore,

$$z_F^0 = \beta(1-e)\big[q\mu_1 + (1-q)\mu_0\big],$$

$$\Delta_F(1) = \beta(1-e)(1-q)(\mu_1 - \mu_0), \qquad \Delta_F(0) = -\beta(1-e)q(\mu_1 - \mu_0).$$

Corollary 2. *Under Assumption 2, if there exists $\beta > 0$ such that $v(d) = -\beta d^2$, then*

$$M_s = \beta(1-e^2)(\mu_s^2 + \tau_s^2), \qquad \tau_s^2 \equiv \mathrm{Var}(W \mid S = s).$$

Therefore,

$$z_F^0 = \beta(1-e^2)\Big[q(\mu_1^2 + \tau_1^2) + (1-q)(\mu_0^2 + \tau_0^2)\Big],$$

$$\Delta_F(1) = \beta(1-e^2)(1-q)\Big[(\mu_1^2 + \tau_1^2) - (\mu_0^2 + \tau_0^2)\Big],$$

$$\Delta_F(0) = -\beta(1-e^2)q\Big[(\mu_1^2 + \tau_1^2) - (\mu_0^2 + \tau_0^2)\Big].$$

3.2.2 WTP for Inspection under Simplified Assumptions

Inspection differs fundamentally from the filter because learning S does not reveal W. Under quasi-linearity, the optimal inspection bid equals the ex-ante value of information.

Lemma 1. *Under Assumption 2, the optimal bid for the inspection is*

$$\tilde{z}^*(\pi) = \tilde{V}_1(\pi) - \tilde{V}_0(\pi), \tag{9}$$

where

$$\tilde{V}_0(\pi) = \max_{a \in \mathcal{A}} \mathbb{E}_\pi \big[-c(a) + v\big(d(W,a)\big) \big]$$

and

$$\tilde{V}_1(\pi) = \mathbb{E}_\pi \left[\max_{a \in \mathcal{A}} \big\{ -c(a) + v\big(d(W,a)\big) \big\} \right].$$

Applying Lemma 1 to the latent-state structure gives the following decomposition.

Proposition 4. *Under Assumption 2, pre-disclosure inspection WTP is*

$$z_I^0 = \tilde{z}^*(\pi^0), \qquad \pi^0 = qF_1 + (1-q)F_0,$$

and post-disclosure inspection WTP after learning $S = s$ is

$$z_I^1(s) = \tilde{z}^*(F_s).$$

Hence the disclosure effect is

$$\Delta_I(s) \equiv z_I^1(s) - z_I^0 = \tilde{z}^*(F_s) - \tilde{z}^*(\pi^0).$$

To sharpen intuition, we impose the same linear mitigation technology state-by-state.

Assumption 3. *The mitigation technology is a binary choice: $\mathcal{A} = \{0,1\}$. There exists $c > 0$ such that $c(a) = ca$, and $d(w,a) = w(1-a)$.*

Assumption 4. *Conditional on $S = s$, beliefs about realized contamination are uniform:*[3]

$$W \mid S = s \sim U(l_s, u_s), \qquad 0 \le l_s < u_s.$$

Corollary 3. *Under Assumptions 2, 3, and 4, if there exists $\beta > 0$ such that $v(d) = -\beta d$, then for any $s \in \{0,1\}$:*

1. *If $u_s \le c/\beta$ or $l_s \ge c/\beta$, then $\tilde{z}^*(F_s) = 0$.*

[3] We impose the uniform distribution here purely to yield clean closed-form thresholds. The main threshold logic holds for any continuous distribution with full support over $[l_s, u_s]$.

2. If $l_s < c/\beta < u_s$, then $\tilde{z}^*(F_s) > 0$.

Corollary 3 formalizes the threshold logic. Inspection is valuable only when the residual uncertainty within a disclosed state straddles the mitigation threshold. If the non-lead state is sufficiently safe that $u_0 \leq c/\beta$, then disclosure pushes households into a corner where further information has no value. If, instead, the lead state satisfies $l_1 < c/\beta < u_1$, then learning $S = 1$ leaves the household in an interior region where direct measurement of W can still affect downstream action, so inspection WTP remains strictly positive. The same logic extends beyond linear damages: diagnostics are valuable when residual within-state uncertainty is action-relevant, not merely when average risk is high.

While our baseline model attributes the value of an inspection to resolving standard risk, the actual household environment may be characterized by profound Knightian uncertainty. Learning $S = 1$ does not merely expose a household to a known variance τ_1^2; it places them in an ambiguous state regarding their exposure to waterborne lead (i.e., specific pipe material, the local water chemistry, corrosion control efficacy, etc). If households are ambiguity averse (Klibanoff et al., 2005), the premium placed on resolving this uncertainty is magnified. In this light, the elicited inspection WTP represents not only the classical value of information but also a strict ambiguity premium triggered by the $S = 1$ disclosure.

This formulation makes clear why filters and inspections need not move together. Filters load on the *between-state* contrast in expected damages, summarized by $M_1 - M_0$. Inspections, by contrast, load on the *within-state* residual uncertainty that remains after the state has been disclosed. A lead disclosure can therefore generate a muted filter response and a strong diagnostic response at the same time.

3.2.3 Predictions

The model delivers three predictions that we take to the data. First, if households begin with overly pessimistic priors about living with a lead line, a non-lead disclosure should reduce filter WTP by correcting that prior. Second, a lead disclosure need not produce a comparably large increase in filter WTP. Proposition 3 shows that the lead-side response is proportional to both $1 - q$ and the gap $M_1 - M_0$. It will therefore be muted when households were already assigning a high probability

to $S = 1$ ex ante and when line material only loosely maps into realized contamination. Third, inspection demand is governed by residual within-state uncertainty rather than by the between-state mean shift alone. If the non-lead state places households in a safe corner, inspection WTP falls after a non-lead disclosure. If the lead state instead leaves households in the region of actionability, inspection WTP rises because direct measurement of W can still change subsequent mitigation choices. The central empirical implication is therefore a portfolio reallocation: non-lead disclosure can reduce demand for both filters and inspections, while lead disclosure can raise demand for precise diagnostics even when the response of a stopgap mitigation good is small.

4 Data and Sample

This section describes the construction of our sampling frame, survey implementation, and the characteristics of the analysis sample that underpins the empirical results. Throughout, we focus on the four cities in which the experiment has already been fielded—Detroit, Indianapolis, Milwaukee, and NYC—and briefly discuss the planned expansion to Minneapolis. We have currently collected 568 valid responses from September 3 to 23, 2025. The completion of the fielding will bring the total target sample to approximately 2,000 respondents across five cities.

4.1 Service Line Records and Study Areas

The starting point for our design is the parcel-level service line inventory data that cities have made public in response to the LCRR. For each of the cities, we obtained the most recent water service line inventory available at the time of experiment implementation and processed the records into GIS layers at the parcel level. Each record reports the material of the water service line connecting the parcel to the water main, typically classified as lead, galvanized, non-lead, or unknown.

Table 1 summarizes these inventories. Across Detroit, Indianapolis, Milwaukee, Minneapolis, and NYC, there are roughly 1.7 million recorded service lines, of which over $385,000 (22\%)$ are identified as lead and $928,000 (53\%)$ as non-lead, with the remainder coded as unknown (including galvanized). These records serve two roles in our project. First, they define the study area by restricting attention to parcels with a known lead status. Second, conditional on successful geospatial matching, they provide the individualized but proxy-based information that is disclosed to treated

respondents.

For the planned extension, we have already processed analogous parcel-level service line data for Minneapolis, which contains a substantial share of lead lines as well. Adding Minneapolis to the experimental sites allows us to test the generalizability of our findings in a fifth, demographically and institutionally distinct city, thereby increasing the external validity of our experiment.

4.2 Current Analysis Sample and Quality Controls

Fieldwork in the four current cities began on September 3, 2025. By September 23, we had obtained 781 complete surveys (394 assigned to control and 387 to treatment), which were trimmed via pre-registered quality screens as shown in Table A2. We first excluded responses from pilot waves, leaving 662 completes. We then dropped respondents whose completion time was below ten minutes and those who failed at least one attention check question. The final analysis sample contains 568 respondents (294 control and 274 treatment). The p-values of the Pearson's Chi-square test between treatment assignment and LSL status shown in Table A2 suggest orthogonality between the two, indicating that response filtering does not compromise the randomized nature of the treatment. Within the cleaned sample, Table 2 reports the cross-tabulation of city, material, and treatment status. In each city, we observe both lead and non-lead homes in both experimental arms, with roughly two-thirds of respondents residing in non-lead buildings and one-third in lead buildings, reflecting the underlying inventories. This joint variation in treatment and true LSL status is central to our empirical strategy.

The sample is well balanced across treatment arms. Table 3 shows that treated and control respondents look similar across a whole host of household characteristics, with all p-values from statistical tests sitting comfortably above conventional significance thresholds. The balance we observe across key household characteristics confirms that our randomization was robust and successful, reinforcing the internal validity of our empirical analysis. Moreover, when we compare our sample to official, population-weighted benchmarks across our study cities, the distributions of key demographics are broadly comparable, with small deviations across average household age and income.

Table 4 compares respondent demographics by recorded service line material to population-weighted benchmarks across our study cities. Households in the lead sample are slightly younger

and have lower incomes than both the non-lead sample and the underlying city populations, and they are substantially less likely to be White (29% vs. 43% in the non-lead group and 36% in the census benchmark). These patterns mirror well-documented environmental justice gradients in lead exposure and related health risks, underscoring why LSL replacement is also an urgent equity concern (Aizer et al., 2018; Hollingsworth and Rudik, 2021). For comparisons in other characteristics by treatment assignment, please refer to Table A3, which includes water usage, perception of health damage from lead, trust in the city's database (and the information received in the treated group), and effectiveness of the Brita pitcher filter. Note that among these variables, only those on water usage are collected before the intervention, while the others are collected at the end of the survey.

We plan to resume the fielding of the survey during February and March 2026 and recruit approximately 300 additional respondents in Minneapolis under the same protocol. Combined with the existing four-city sample, this will yield a five-city experiment with about 2,000 respondents, improving statistical power and allowing us to probe heterogeneity in treatment effects across distinct regulatory and demographic environments.

5 Results

5.1 Pre-intervention Baseline Analytics

We begin by documenting the pre-intervention characteristics of the experimental sample and verifying that randomization produced comparable treatment and control groups. Tables 5 and 6 report pre-treatment beliefs and WTPs overall and by material. Prior to information disclosure, treatment and control groups look similar: respondents assign on average a probability of $\sim 57\%$ that their home has an LSL, expect ~ 25 tap-water samples to exceed the lead standard, and express mean WTPs for a pitcher filter and for an inspection in the $\$65 - \$75$ range. Differences between arms are small and statistically insignificant within both the lead and non-lead strata. This supports our use of post-minus-pre changes as the main outcome in the empirical specifications.

At the same time, the levels of these priors are striking: in the underlying inventories, only about 29% of homes in our analysis sample are actually recorded as having LSLs, making average subjective priors almost twice as high as the objective LSL share. The average prior on lead

exceedance in tap water is also over twice the EPA safety limit (10 samples). This large ex ante gap between perceived and actual LSL prevalence underscores both the salience of waterborne lead in households' minds and the potential value of inventory-based disclosure in correcting systematically pessimistic priors about service-line risk. It also contrasts with recent evidence from highly polluted ambient-air settings, where households often underestimate actual pollution levels or the returns to mitigation (Ahmad et al. (2022); Chowdhury et al. (2025); Greenstone et al. (2021)).

5.2 Main Results

We now turn to the effects of disclosing building-specific LSL status on beliefs and WTPs. Figures 3-5 illustrate the changes in our outcome variables caused by the intervention, broken out by experimental arm and material type; Table 7 presents ITT estimates from Equation (1). We organize the results in three parts: belief updating (Columns $1-3$), mitigation demand (Column 4), and diagnostic demand (Column 5).

Beliefs about LSL presence respond strongly and in the direction implied by the disclosed information. Figure 3a plots the change in the percentage belief reported before and after the information intervention; that is, $P_i^{post} - P_i^{pre}$. The figure indicates that information disclosure increases the perceived likelihood among lead respondents and decreases it among non-lead respondents. Based on Column (1) of Table 7, among respondents whose addresses are recorded as having lead service lines, treatment increases the perceived probability of having an LSL by about 21 percentage points relative to non-lead respondents. This effect corresponds to roughly 0.8 pre-treatment standard deviation of LSL beliefs. Panels (a) and (b) of Figure 6 illustrate these effects by reporting the estimated treatment and interaction coefficients along with their 95% confidence intervals.

In addition, disclosure shifts beliefs about tap-water contamination. Figure 4a reveals a divergence in the treated group: respondents with LSLs increase their estimates of lead exceedances, while those without LSLs lower their estimates. Meanwhile, the control group shows little change. Column (2) of Table 7 shows that, for lead addresses, the information intervention raises the perceived number of lead exceedances (out of 100 hypothetical samples) by about 6.7 samples relative to non-lead addresses; this effect is economically meaningful, though estimated with substantial uncertainty in the current sample. Figure 6 visualize the corresponding treatment and interaction estimates with 95% confidence intervals in Panels (c) and (d). Column (4) of Table 11 shows a

more precisely estimated effect under the learning-rate specification in Equation (3), of similar magnitude (8.84).

Information disclosure increases belief confidence. In Figure 3b, we treat P_i^{pre} and P_i^{post} as the probability parameters of Bernoulli distributions to calculate the pre- and post-standard deviations. For both the lead respondents and the non-lead respondents, we observe a larger decrease in the standard deviation among treated respondents relative to those in the control arm. Specifically, the treated non-lead group appears to exhibit the most significant reduction in uncertainty. Figure 4b plots the change in confidence regarding the number of lead exceedances in 100 kitchen tap water samples. The figure confirms an increase in confidence among the treated respondents compared with the control group.

Column (3) of Table 7 shows that, pooling across lead and non-lead homes, assignment to treatment increases respondents' confidence in their beliefs by about 0.22 point on a 1–5 scale, approximately 0.23 standard deviation of the baseline confidence distribution. This result is illustrated in Figure 7, which shows the point estimate and the 95% confidence interval for Treatment. This pattern is consistent with disclosure narrowing the perceived distribution of possible states, even when the mean risk moves up or down. The joint movement of beliefs and confidence also resonates with the broader literature on self-confidence and information acquisition, which shows that individuals adjust confidence in response to signals in ways that shape subsequent choices and welfare (Möbius et al., 2022). While our context is simpler and does not model strategic confidence management, the pattern that confidence responds systematically to inventory signals—alongside beliefs—underscores that disclosure changes not only perceived risk levels but also how certain households feel about those beliefs.

Turning to defensive behavior, we find clear evidence that providing proxy-based information changes the demand for diagnostics, while the effect on mitigation is not precisely estimated for lead respondents. Figure 5a plots the distribution of the change in WTP for the Brita pitcher-style filter; it suggests a treatment-induced decrease among the non-lead group but statistically imprecise effect for the lead group. In contrast, Figure 5b reveals a sharp treatment-induced divergence in demand for diagnostics: WTP increases for the treated lead group and decreases for the treated non-lead group, while the control group remains relatively stable.

Column (4) of Table 7 confirms that, among non-lead respondents, the treatment reduces WTP

for a pitcher-style filter by about \$8 on average (roughly 0.18 baseline standard deviation of the non-lead respondents). For lead respondents, the estimate of the treatment effect (i.e., the sum of the coefficients on Treatment and Lead $\times$ Treatment) is less precisely estimated (point estimate $= -\$2.30$, 95% CI $[-\$10.19,\ \$5.59]$). We do not interpret this near-zero estimate as evidence that LSL households fail to update. Rather, taken together with the belief results and the strong increase in inspection WTP, the pattern is more consistent with households distinguishing between learning that they are attached to a risky infrastructure type and learning that their own tap water is contaminated. A lead-line disclosure is bad news about an upstream hazard, but it remains an imperfect proxy for realized exposure at the tap, which also depends on corrosion control, premise plumbing, and household water use. Under this interpretation, lead disclosure need not generate a large increase in demand for a generic mitigation good such as a pitcher filter, because households may prefer to resolve the remaining household-specific uncertainty first.

Conversely, Column (5) of Table 7 shows a significant positive treatment–lead interaction: among lead-address households, the information intervention increases WTP for a professional kitchen-tap inspection by about \$16 relative to non-lead households, corresponding to roughly 0.36 of the baseline standard deviation. Specifically, the LSL group increased their WTP by \$8.91, whereas the non-LSL group decreased theirs by \$7.26. Figure 8 visualizes the corresponding treatment and interaction estimates for inspection WTP in Panels (c) and (d). This asymmetry between filters and inspections is consistent with households treating the inspection as the more targeted good. Non-lead disclosure resolves a substantial amount of pessimistic uncertainty and lowers the value of both generic mitigation and further information. Lead disclosure, by contrast, heightens concern about the underlying infrastructure while leaving the exact contamination level in the household's drinking water unresolved, thereby raising the value of direct diagnostics even when the valuation of a stopgap filter does not move much.

We interpret these results through the model as evidence consistent with a threshold-based value-of-information mechanism, augmented by the institutional fact that service-line material is only a proxy for realized exposure. For LSL households, disclosure appears to raise perceived infrastructure risk while also reducing uncertainty. In the model, this combination need not raise inspection WTP in general. It does so only when households remain in a region where additional information could still affect subsequent mitigation choices. Our estimates are therefore most

naturally read as consistent with lead-household beliefs moving into, or remaining within, such an interior region, so that the increase in perceived risk is not dominated by the reduction in uncertainty. At the same time, the muted filter response suggests that respondents do not treat the inventory record as sufficient to infer severe realized contamination at the tap. This interpretation is deliberately local: the model does not imply that higher mean risk generically raises diagnostic demand, only that it can do so when information remains action-relevant.

Table 9 assesses whether randomly switching the order in which the two WTP questions are asked affects respondents' stated values. Column (1) shows that when respondents are first asked how much they are willing to give up their bonus to obtain a professional tap water lead inspection—before being asked about a lead-filtering pitcher—their stated WTP for the filter is about $8 lower on average, and this difference is statistically significant. In contrast, column (2) shows that the question order does not affect WTP for professional inspection.

This pattern highlights the importance of our research in studying demand for more accurate exposure assessment when individuals are provided with proxy information. Specifically, it suggests that presenting the inspection option first may crowd out respondents' valuation of immediate mitigation options. Columns (3)–(6) further show that switching the question order does not affect changes in WTP or posterior WTP once prior WTP is controlled for, indicating that order effects do not meaningfully affect our main results.

5.3 Robustness Checks

We assess the robustness of these findings across several pre-specified dimensions in our analysis plan. First, we augment equation (1) with baseline covariates capturing household income and the presence of children aged 10 or younger. These variables are potentially important both for the perceived stakes of lead exposure and for access to mitigation technologies. The coefficients on the treatment–lead interaction remain practically unchanged when these controls are included (see Table 8).

Second, we estimate a "post-on-pre" version of the main specification, as specified in Equation (2). The resulting estimates, summarized in Table 10, confirm that our conclusions are not an artifact of using change scores; the implied treatment–lead effects on posterior beliefs and WTPs are of similar order of magnitude to the baseline ITT estimates.

Third, we exploit the "perception gap" structure of our design and re-estimate the main specification by replacing the lead-status indicator with the pre-treatment perception gap, as specified in Equation (3). This alternative parametrization follows the learning literature and allows us to relate belief updating directly to the size and sign of prior errors. Table 11 reports coefficient estimates of 2.20 and 8.84 for perception gap $\times$ Treatment in Columns (2) and (4). This means that the perception-gap specifications yield patterns that closely resemble Bayesian updating: larger prior errors are associated with proportionally larger revisions in the direction of the disclosed inventory record. This structure echoes Armantier et al. (2016), who show that inflation expectations update in proportion to prior forecast errors when households receive price information. In our context, the fact that the magnitude of belief changes scales with the perception gap, rather than being a constant shift across respondents, helps rule out interpretations based purely on generic salience or experimenter demand. Instead, it suggests that respondents treat LSL inventory information as a noisy signal and incorporate it in a reasonably disciplined way.

5.4 Heterogeneity Analysis

This section outlines the heterogeneity analyses we plan to carry out once we complete the fielding. Our experimental design provides several natural dimensions of heterogeneity. First, because the sample spans four distinct cities (with Minneapolis to be added), we can ask whether the impact of disclosure differs across water systems with different LSL prevalence, regulatory histories, and communication practices. In the completed paper, we will present estimates of equation (1) allowing the treatment–lead interaction to vary by city and will relate these differences to features of local implementation.

Second, we will explore heterogeneity by baseline vulnerability and information environment. This includes interactions with household income, the presence of young children, education, race/ethnicity, baseline beliefs, and confidence. These analyses will shed light on whether more disadvantaged or more uncertain households react more strongly to individualized proxy information, and whether disclosure narrows or widens existing gaps in defensive investments.

Third, we will examine whether the shift from filters to inspections among lead households is concentrated among those with significant ex ante perception gaps, as the model suggests. Estimating equation (1) with interactions between treatment assignment and terciles of the perception-gap

distribution will allow us to test whether households who were initially most misinformed reallocate defensive resources more aggressively when confronted with building-specific LSL records.

The combination of these heterogeneity analyses will enable us to speak not only to the average effects of LSL disclosure but also to its distributional consequences and the potential for targeted information or subsidies to enhance both efficiency and equity in future LSL replacement programs.

5.5 Welfare Calculations and Policy Implications

Our results bear directly on LSL replacement policy design. The 2021 Infrastructure Investment and Jobs Act allocates $15 billion for LSL replacement, while LCRR mandates inventory disclosure but remains largely silent on complementary policies. We discuss three implications.

First, inventory disclosure generates welfare gains through correcting the WTP of mis-priced goods. Prior to the intervention, non-lead households held diffuse, unrealistically pessimistic priors that induced a positive valuation for a mitigation good they did not technically need. The information shock compresses this deadweight loss by reducing wasteful filter expenditure by about $8 per non-lead household and reducing inspection WTP by about $7 per non-lead household. Applied to roughly $928,000$ non-lead service lines in our five-city study area, these two channels together correspond to avoided inefficient defensive expenditures on the order of $13.9 million. On the diagnostic margin, lead households reveal an average post-intervention increase of about $9 in WTP for a professional inspection, implying that subsidized testing could unlock roughly $3.5 million in latent demand across the approximately $385,000$ lead parcels. Taken together, these three channels alone suggest a minimum static social value of $\sim$ $17 million in our sample cities alone, even before accounting for health improvements, permanent pipe replacement, or general equilibrium effects.

Second, liquidity constraints likely bind for vulnerable households. Table 4 shows lead households are lower-income and disproportionately non-White, consistent with environmental justice literature (Banzhaf et al., 2019). Despite these tighter budgets, lead households exhibit high inspection WTP ($\sim$ $76) when provided experimental endowments of $150. This suggests the barrier to seek the service of exposure investigation might not be low evaluation of the service but the budget constraint. Market-based solutions that rely on households purchasing their own tests post-disclosure may therefore be regressive. An optimal policy should couple inventory disclosure with subsidized verification for identified lead households.

Third, disclosure should be paired with salient communication. Our intervention provides crisp, individualized information attributed to official city records. Real-world LCRR implementation varies widely in salience: some utilities mail postcards; others bury inventories on an online search portal. Our results suggest disclosure matters but only if households actually receive and attend to the information.

6 Conclusion

This paper examines how proxy-based, individualized environmental information shapes beliefs and defensive investments in the context of waterborne lead exposure. Leveraging parcel-level service line inventories that U.S. cities assembled under EPA's Lead and Copper Rule Revision, we design and implement a pre-registered, within-city randomized experiment in five large urban systems. The intervention discloses respondents' building-specific LSL status and, before and after disclosure, elicits beliefs about having an LSL and tap-water exceedances, confidence, and incentive-compatible WTP for a pitcher filter and for a professional kitchen-tap inspection. A stylized model distinguishes uncertainty over latent pipe material from uncertainty over realized tap-water contamination conditional on pipe material and links these two layers of uncertainty to demand for stopgap mitigation and precise diagnostics.

Three main empirical findings emerge. First, individualized disclosure induces sharp belief updating: among homes recorded as having an LSL, the perceived probability of having an LSL rises substantially relative to non-lead homes, and belief confidence increases across the treated sample. These patterns are consistent with Bayesian learning from a high-signal proxy, in line with evidence on environmental information and belief revision in other settings (Neidell, 2010; Barwick et al., 2024; Conell-Price and Mulder, 2024). Second, disclosure reallocates defensive effort away from generic mitigation and toward diagnostics. On average, filter WTP falls for non-lead households, while the lead-household filter response is statistically imprecise and centered near zero; at the same time, inspection WTP rises sharply for lead-address households and falls for non-lead households. We interpret this pattern as evidence that households distinguish between infrastructure risk and realized household exposure. Non-lead disclosure corrects overly pessimistic priors and reduces the value of both stopgap mitigation and further information. Lead disclosure, by contrast, signals

that the underlying infrastructure is riskier without fully resolving contamination at the tap, which makes direct diagnostics more attractive than a generic protective good. Third, these patterns are robust to alternative specifications, including post-on-pre formulations and updating-based parametrizations using the perception gap between priors and inventory records. Together, we conclude that the LSL disclosure generates a minimum social value of $\sim$ \$17 million in our sample cities.

The model developed in Section 3 helps organize these findings by distinguishing uncertainty over infrastructure type from uncertainty over realized household exposure. Households begin with a subjective prior over a latent state $S \in \{0, 1\}$, where $S = 1$ denotes a lead service line, and realized tap-water contamination is then drawn from a state-contingent distribution F_s. Disclosure reveals S, but not the contamination draw itself. In this setting, filter WTP is governed by the between-state contrast in expected damages. When households start from pessimistic priors about living with a lead line, a non-lead disclosure can sharply reduce filter demand, while a lead disclosure need not raise it much if line material is only a noisy predictor of realized contamination. Inspection WTP, by contrast, is the ex-ante value of learning realized contamination after the proxy state has already been revealed. It therefore depends on residual within-state uncertainty: it is low when disclosure places households in a safe corner where subsequent action is obvious, and high when disclosure leaves them in a region where direct measurement can still change mitigation choices. The empirical pattern in our data—a strong negative filter response for non-lead households, a muted filter response for lead households, and a pronounced rise in inspection WTP for lead households—is consistent with exactly this combination of pessimistic priors, noisy proxy-to-outcome mapping, and residual action-relevant uncertainty. Taken together with the perception-gap regressions, these patterns place our findings squarely within a broader body of work documenting disciplined belief updating in response to information (e.g. Armantier et al., 2016). They also highlight that confidence is an elastic object influenced by signals in ways reminiscent of experimental evidence on self-confidence management (Möbius et al., 2022).

Normatively, our results bear directly on current policy debates. The United States is poised to spend tens of billions of dollars over the coming decade to identify and replace LSLs, combining federal infrastructure funding with state and local programs (USEPA, 2023). Knowing where lead pipes are is a necessary input into any efficient replacement strategy, but it is not sufficient: the

welfare returns to these investments depend on whether households understand their risk and adjust behavior accordingly. Our experiment shows that inventory-based disclosure can realign beliefs with the underlying infrastructure and shift demand toward individualized diagnostics when risk is elevated, while discouraging low-return mitigation when risk is low. These features are desirable from a welfare perspective if households internalize the health benefits of accurate exposure assessment and subsequent action, and they complement the large body of evidence documenting the long-run harms of lead exposure for health and human capital (e.g., Aizer et al., 2018; Clay et al., 2014; Dave and Yang, 2022; Hollingsworth et al., 2022; Marcus, 2023). From a policy design standpoint, our findings suggest that pairing inventory disclosure with subsidized testing, targeted outreach in lead-dense neighborhoods, and ultimately subsidized replacement could yield more cost-effective and equitable reductions in exposure than either information or subsidies alone.

In terms of external validity, our findings likely generalize to other settings where regulators disclose infrastructure risk factors rather than direct pollution measurements (e.g., flood zone designations and proximity to Superfund sites). In each case, households receive a proxy signal and must decide whether to invest in diagnostics or stopgap mitigation. Our finding that proxy disclosure triggers demand for verification suggests information policy design should anticipate this complementarity. Pairing disclosure with subsidized diagnostics may be more cost-effective than disclosure alone.

Finally, our results demonstrate a critical environmental justice dimension of the LCRR. We document that households with lead service lines are systematically lower-income and less likely to be White (see Table 4). Despite these tighter budget constraints, these households exhibited a high willingness to pay for professional inspections once provided with the inventory data and the liquidity of the experimental endowment. This suggests that the primary barrier to exposure assessment in vulnerable communities may not be a lack of valuation or awareness, but a liquidity constraint. Market-based solutions that rely on households purchasing their own tests after disclosure may therefore be regressive. Optimal policy design should likely couple inventory disclosure with subsidized verification for identified "lead" households to close the gap between the desire for safety and the ability to pay.

Several open questions remain and point to a broader research agenda on waterborne lead and information. It is important to note that our $17 million estimate is a lower bound. It excludes

(1) health benefits from improved mitigation and eventual pipe replacement, (2) option value of information for future decisions, (3) spillovers to untreated neighbors, and (4) general equilibrium effects on property values. Back-of-the-envelope calculations using epidemiological dose-response functions (e.g., Hollingsworth and Rudik, 2021; Keiser et al., 2023) suggest the health benefits of targeted replacement could dwarf our measured WTP changes. However, our experiment focuses on short-run responses to a one-time disclosure. Linking public disclosure and experimentally elicited WTPs to administrative data on actual tap-water testing, filter adoption, or enrollment in LSL replacement programs would tighten the connection between stated and revealed preferences and sharpen the welfare calculus. Understanding how beliefs, confidence, and defensive investments evolve as households receive repeated inventory notifications, local outreach, and program offers is also an important topic for future work. Our results provide one step toward filling this gap and illustrate how experimental variation in disclosure guided by theory can be used to evaluate and improve information policies in drinking-water regulation.

References

Ahmad, Husnain, Matthew Gibson, Fatiq Nadeem, Sanval Nasim, and Arman Razaee. 2022. "Forecasts: Consumption, Production, and Behavioral Responses."

Aizer, Anna, and Janet Currie. 2019. "Lead and Juvenile Delinquency: New Evidence from Linked Birth, School, and Juvenile Detention Records." *The Review of Economics and Statistics* 101 (4): 575–587. 10.1162/rest_a_00814.

Aizer, Anna, Janet Currie, Peter Simon, and Patrick Vivier. 2018. "Do Low Levels of Blood Lead Reduce Children's Future Test Scores?" *American Economic Journal: Applied Economics* 10 (1): 307–341. 10.1257/app.20160404.

Allcott, Hunt. 2011. "Social norms and energy conservation." *Journal of public Economics* 95 (9-10): 1082–1095.

Allcott, Hunt, and Todd Rogers. 2014. "The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation." *American Economic Review* 104 (10): 3003–3037.

Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert Van der Klaauw, and Basit Zafar. 2016. "The price is right: Updating inflation expectations in a randomized price information experiment." *Review of Economics and Statistics* 98 (3): 503–523.

Bakkensen, Laura A, and Lala Ma. 2020. "Sorting over flood risk and implications for policy reform." *Journal of Environmental Economics and Management* 104 102362.

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

Barwick, Panle Jia, Shanjun Li, Liguo Lin, and Eric Yongchen Zou. 2024. "From Fog to Smog: The Value of Pollution Information." *American Economic Review* 114 (5): 1338–1381. 10.1257/aer.20200956.

Becker, Gordon M., Morris H. Degroot, and Jacob Marschak. 1964. "Measuring utility by a single-response sequential method." *Behavioral Science* 9 (3): 226–232. 10.1002/bs.3830090304.

Bergemann, Dirk, and Stephen Morris. 2019. "Information design: A unified perspective." *Journal of Economic Literature* 57 (1): 44–95.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2022. "Overreaction and diagnostic expectations in macroeconomics." *Journal of Economic Perspectives* 36 (3): 223–244.

Bui, Linda TM, Ron Shadbegian, Heather Klemick, Dennis Guignet, Rebecca Margolit, and Anh Hoang. 2024. "How Effective Are Secondary Interventions at Improving Health Outcomes in Children Exposed to Lead in Early Childhood?" *Review of environmental economics and policy* 18 (2): 261–278.

Bursztyn, Leonardo, Georgy Egorov, Ingar Haaland, Aakaash Rao, and Christopher Roth. 2023. "Justifying dissent." *The Quarterly Journal of Economics* 138 (3): 1403–1451.

Caplin, Andrew, and Mark Dean. 2015. "Revealed preference, rational inattention, and costly information acquisition." *American Economic Review* 105 (7): 2183–2203.

Chowdhury, Ashfaqul, Teevrat Garg, Maulik Jagnani, and Martin Mattsson. 2025. "Misbeliefs, Experience, and Technology Adoption: Evidence from Air Purifiers in Bangladesh." *Available at SSRN 5284948*.

Clay, Karen, Edson R Severnini, and Xiao Wang. 2025. "The Hidden Toll of Airborne Lead: Infant Mortality Impacts of Industrial Lead Pollution." *Working paper*.

Clay, Karen, Werner Troesken, and Michael Haines. 2014. "Lead and Mortality." *The Review of Economics and Statistics* 96 (3): 458–470. 10.1162/REST_a_00396.

Conell-Price, Lynn, and Philip Mulder. 2024. "Does Good News Wash out the Bad? Belief Updating and Climate Risk Information." March. 10.2139/ssrn.4811717.

Dave, Dhaval M., and Muzhe Yang. 2022. "Lead in drinking water and birth outcomes: A tale of two water treatment plants." *Journal of Health Economics* 84 102644. 10.1016/j.jhealeco.2022.102644.

Dupas, Pascaline. 2011. "Do teenagers respond to HIV risk information? Evidence from a field experiment in Kenya." *American Economic Journal: Applied Economics* 3 (1): 1–34.

Dupas, Pascaline, Basimenye Nhlema, Zachary Wagner, Aaron Wolf, and Emily Wroe. 2023. "Expanding access to clean water for the rural poor: experimental evidence from Malawi." *American Economic Journal: Economic Policy* 15 (1): 272–305.

Ferrie, Joseph P., Karen Rolf, and Werner Troesken. 2012. "Cognitive disparities, lead plumbing, and water chemistry: Prior exposure to water-borne lead and intelligence test scores among World War Two U.S. Army enlistees." *Economics & Human Biology* 10 (1): 98–111. 10.1016/j.ehb.2011.09.003.

Gao, Xuwen, Ran Song, and Christopher Timmins. 2023. "Information, migration, and the value of clean air." *Journal of Development Economics* 163 103079.

Garza, Anibal, Rosario Vega, and Enrique Soto. 2006. "Cellular mechanisms of lead neurotoxicity." *Medical science monitor* 12 (3): RA57.

Greenstone, Michael, Kenneth Lee, and Harshil Sahai. 2021. "Indoor air quality, information, and socioeconomic status: Evidence from Delhi." In *AEA Papers and Proceedings*, Volume 111. 420–424, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Grönqvist, Hans, J Peter Nilsson, and Per-Olof Robling. 2020. "Understanding how low levels of early lead exposure affect children's life trajectories." *Journal of Political Economy* 128 (9): 3376–3433.

Grossman, Daniel S., and David J.G. Slusky. 2019. "The Impact of the Flint Water Crisis on Fertility." *Demography* 56 (6): 2005–2031. 10.1007/s13524-019-00831-0.

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart. 2023. "Designing Information Provision Experiments." *Journal of Economic Literature* 61 (1): 3–40. 10.1257/jel.20211658.

Hausman, Catherine, and Samuel Stolper. 2021. "Inequality, information failures, and air pollution." *Journal of Environmental Economics and Management* 110 102552.

Hollingsworth, Alex, Jiafang Mike Huang, Ivan Rudik, and Nicholas J. Sanders. 2022. "A Thousand Cuts: Cumulative Lead Exposure Reduces Academic Achievement." *Journal of Human Resources*. 10.3368/jhr.0222-12169R2, Publisher: University of Wisconsin Press Section: Articles.

Hollingsworth, Alex, and Ivan Rudik. 2021. "The Effect of Leaded Gasoline on Elderly Mortality: Evidence from Regulatory Exemptions." *American Economic Journal: Economic Policy* 13 (3): 345–373. 10.1257/pol.20190654.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge is (less) power: Experimental evidence from residential energy use." *American Economic Review* 104 (4): 1417–1438.

Keiser, David A, Bhash Mazumder, David Molitor, and Joseph S Shapiro. 2023. "Water works: Causes and consequences of safe drinking water in America." In *National Center for Environmental Economics Seminar, Washington, DC, April*, Volume 13. 853–872.

Keiser, David A, and Joseph S Shapiro. 2019. "US water pollution regulation over the past half century: burning waters to crystal springs?" *Journal of Economic Perspectives* 33 (4): 51–75.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. "A smooth model of decision making under ambiguity." *Econometrica* 73 (6): 1849–1892.

Litman, Leib, and Jonathan Robinson. 2020. *Conducting online research on Amazon Mechanical Turk and beyond*. Sage Publications.

Ma, Lala, Margaret Walls, Matthew Wibbenmeyer, and Connor Lennon. 2024. "Risk disclosure and home prices: Evidence from California wildfire hazard zones." *Land Economics* 100 (1): 6–21.

Madajewicz, Malgosia, Alexander Pfaff, Alexander Van Geen, Joseph Graziano, Iftikhar Hussein, Hasina Momotaj, Roksana Sylvi, and Habibul Ahsan. 2007. "Can information alone change behavior? Response to arsenic contamination of groundwater in Bangladesh." *Journal of development Economics* 84 (2): 731–754.

Marcus, Michelle. 2023. "Burying the Lead: Effects of Public Lead Service Line Replacements on Blood Lead Levels and Property Values." July, https://michellemmarcus.com/wp-content/uploads/2023/07/marcus_buryingthelead_072223.pdf.

Menke, Andy, Paul Muntner, Vecihi Batuman, Ellen K Silbergeld, and Eliseo Guallar. 2006. "Blood lead below 0.48 μmol/L (10 μg/dL) and mortality among US adults." *Circulation* 114 (13): 1388–1394.

Metcalfe, Robert D., and Sefi Roth. 2025. "Making the Invisible Visible: The Impact of Revealing Indoor Air Pollution on Behavior and Welfare." February. 10.3386/w33510.

Möbius, Markus M, Muriel Niederle, Paul Niehaus, and Tanya S Rosenblat. 2022. "Managing self-confidence: Theory and experimental evidence." *Management Science* 68 (11): 7793–7817.

Neidell, Matthew. 2010. "Air quality warnings and outdoor activities: evidence from Southern California using a regression discontinuity design." *Journal of Epidemiology & Community Health* 64 (10): 921–926. 10.1136/jech.2008.081489.

Reyes, Jessica Wolpaw. 2015. "Lead Exposure and Behavior: Effects on Antisocial and Risky Behavior Among Children and Adolescents." *Economic Inquiry* 53 (3): 1580–1605. 10.1111/ecin.12202.

Stigler, George J. 1961. "The economics of information." *Journal of political economy* 69 (3): 213–225.

USEPA. 2020. "Revised Lead and Copper Rule." December, https://www.epa.gov/ground-water-and-drinking-water/revised-lead-and-copper-rule.

USEPA. 2023. "EPA's 7th Drinking Water Infrastructure Needs Survey and Assessment." August, https://www.epa.gov/dwsrf/epas-7th-drinking-water-infrastructure-needs-survey-and-assessment, Archive Location: United States.

Zivin, Joshua Graff, and Matthew Neidell. 2009. "Days of haze: Environmental information disclosure and intertemporal avoidance behavior." *Journal of Environmental Economics and Management* 58 (2): 119–128.

7 Figures and Tables

Figures

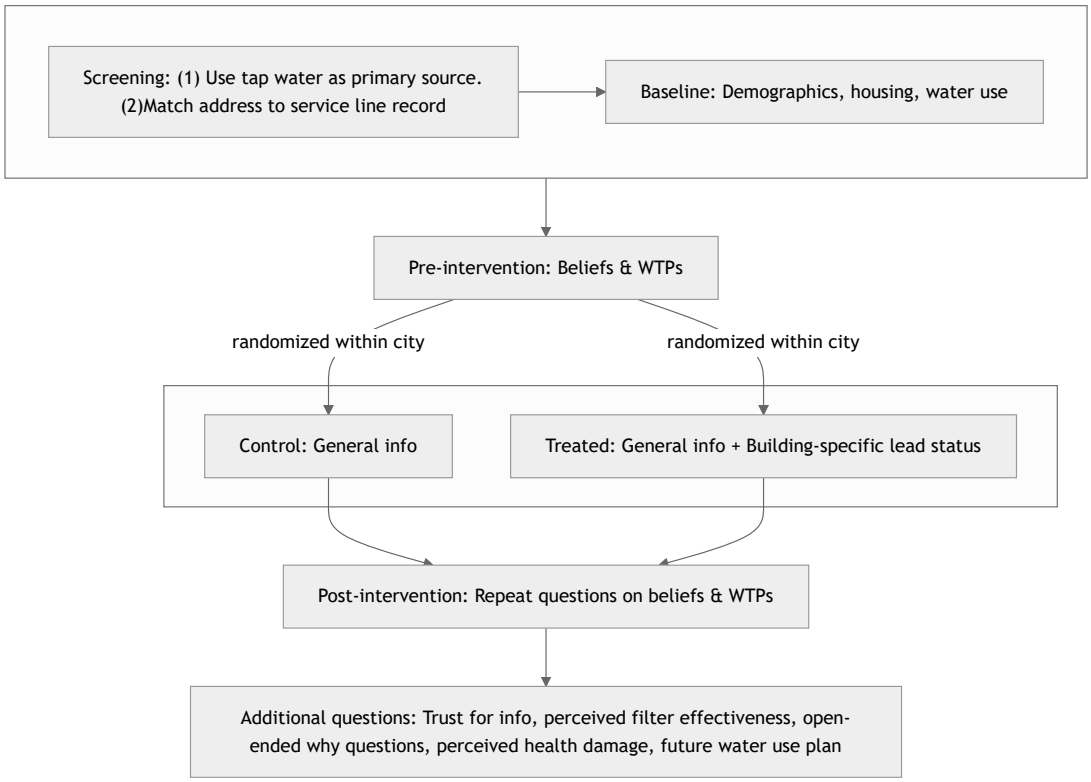


Figure 1: Survey Flow of the Experiment

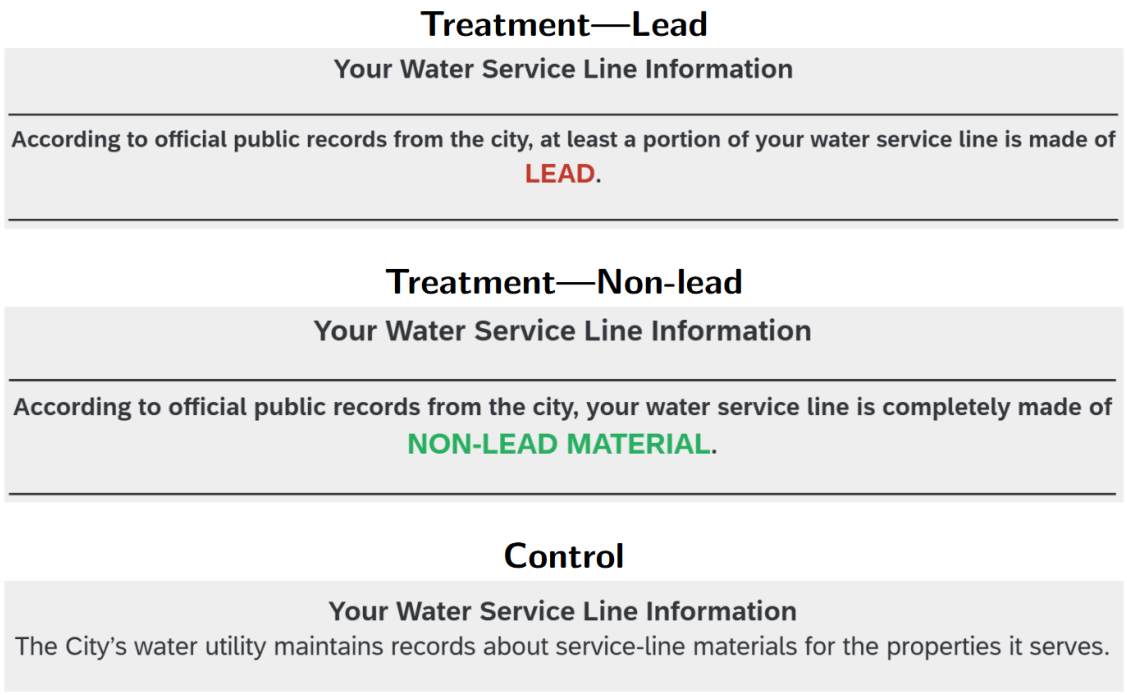


Figure 2: Screenshots of Intervention Messages

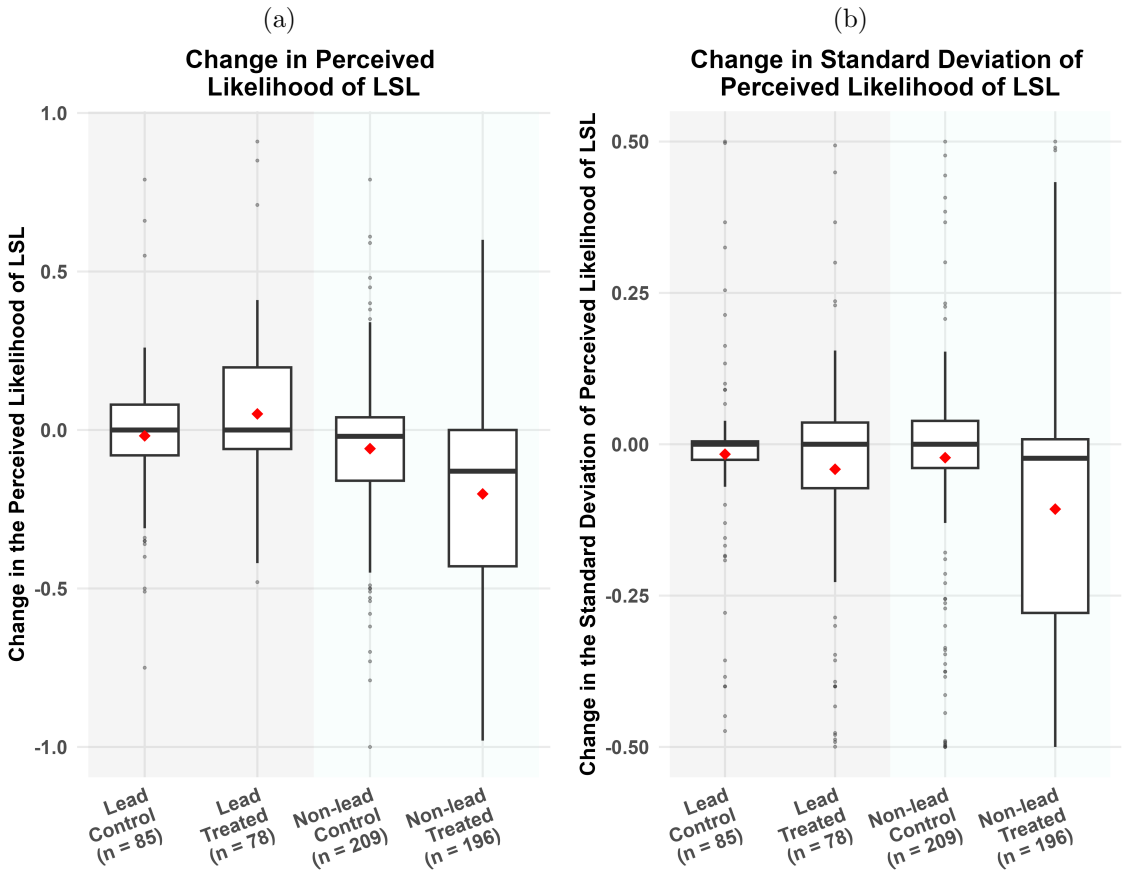


Figure 3: Change in the perceived likelihood of LSL presence among each group. For each box-and-whisker plot, the bottom of the box, the thick line, and the top of the box denotes the 25th, the 50th, and 75th percentile of the distribution, respectively. The red marker indicates the mean value. Panel (a) shows the result for the change in the percentage point belief reported before and after the information intervention, $P_i^{post} - P_i^{pre}$. In Panel (b), we treat P_i^{pre} and P_i^{post} as the probability of Bernoulli distribution and calculate the pre- and post-standard deviation respectively and then calculate the change.

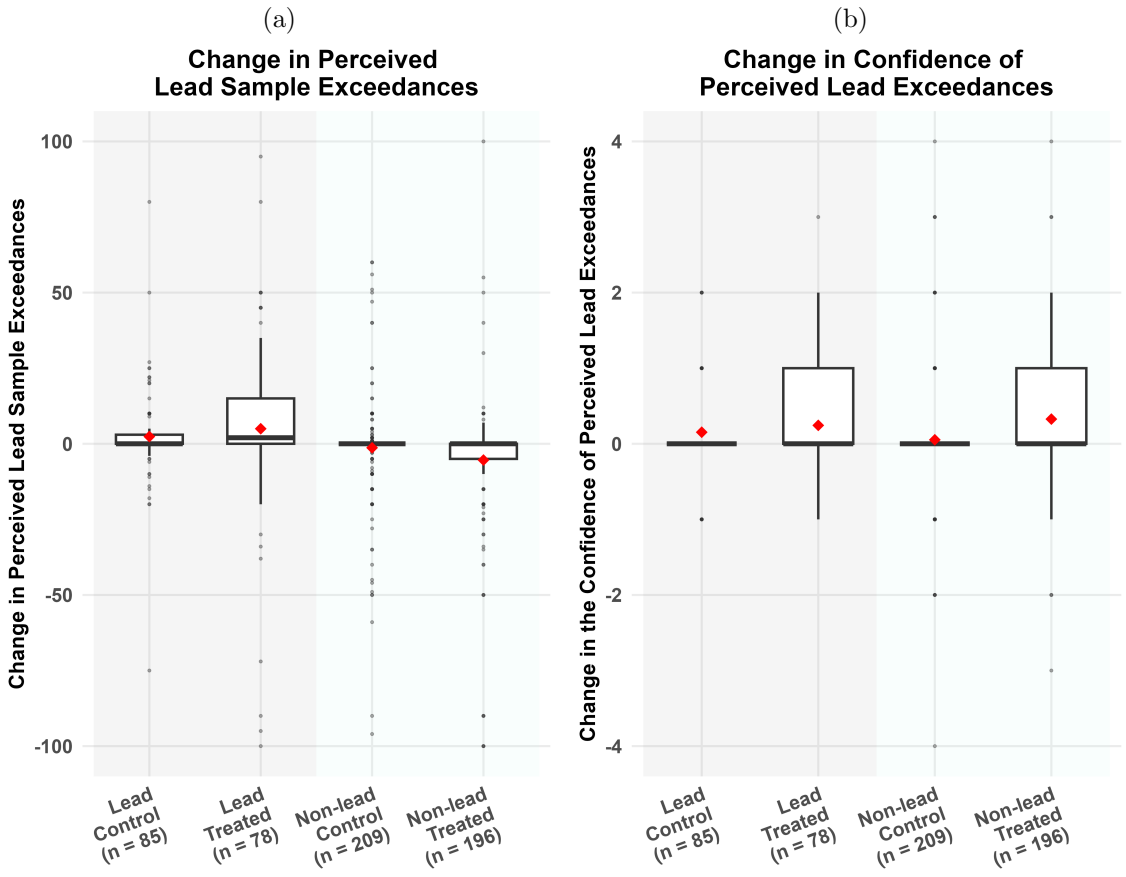


Figure 4: Belief change in lead sample exceedances among each group. For each box-and-whisker plot, the bottom of the box, the thick line, and the top of the box denote the 25th, 50th, and 75th percentiles of the distribution, respectively. The red marker indicates the mean value. Panel (a) shows the change of perceived number of water samples that exceed the federal limit of 0.01mg/L in 100 kitchen water samples. Panel (b) shows the change in the confidence of that belief measured on a 1-5 scale.

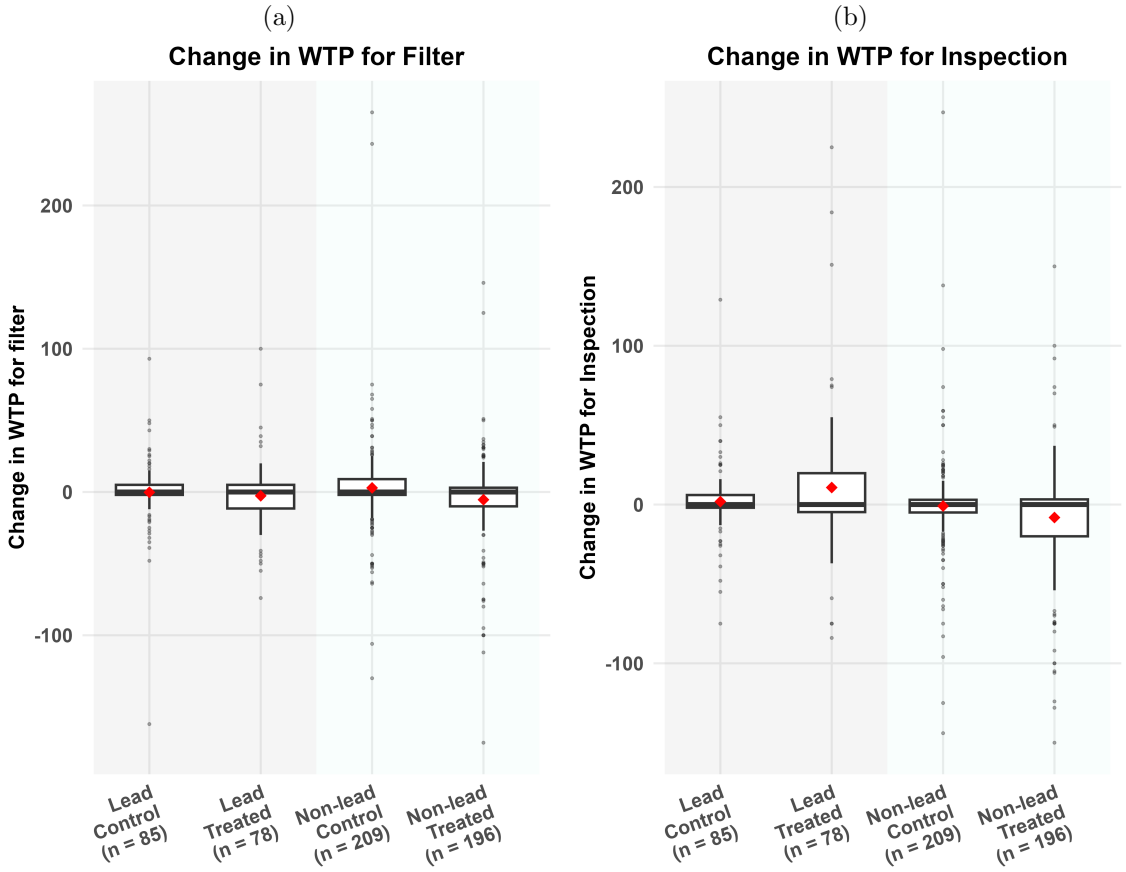


Figure 5: WTP changes among each group. For each box-and-whisker plot, the bottom of the box, the thick line, and the top of the box denote the 25th, the 50th, and the 75th percentile of the distribution, respectively. The red marker indicates the mean value. Panel (a) shows the change in the WTP for the pitcher-style filter. Panel (b) shows the change in the WTP for the professional tap water lead inspection.

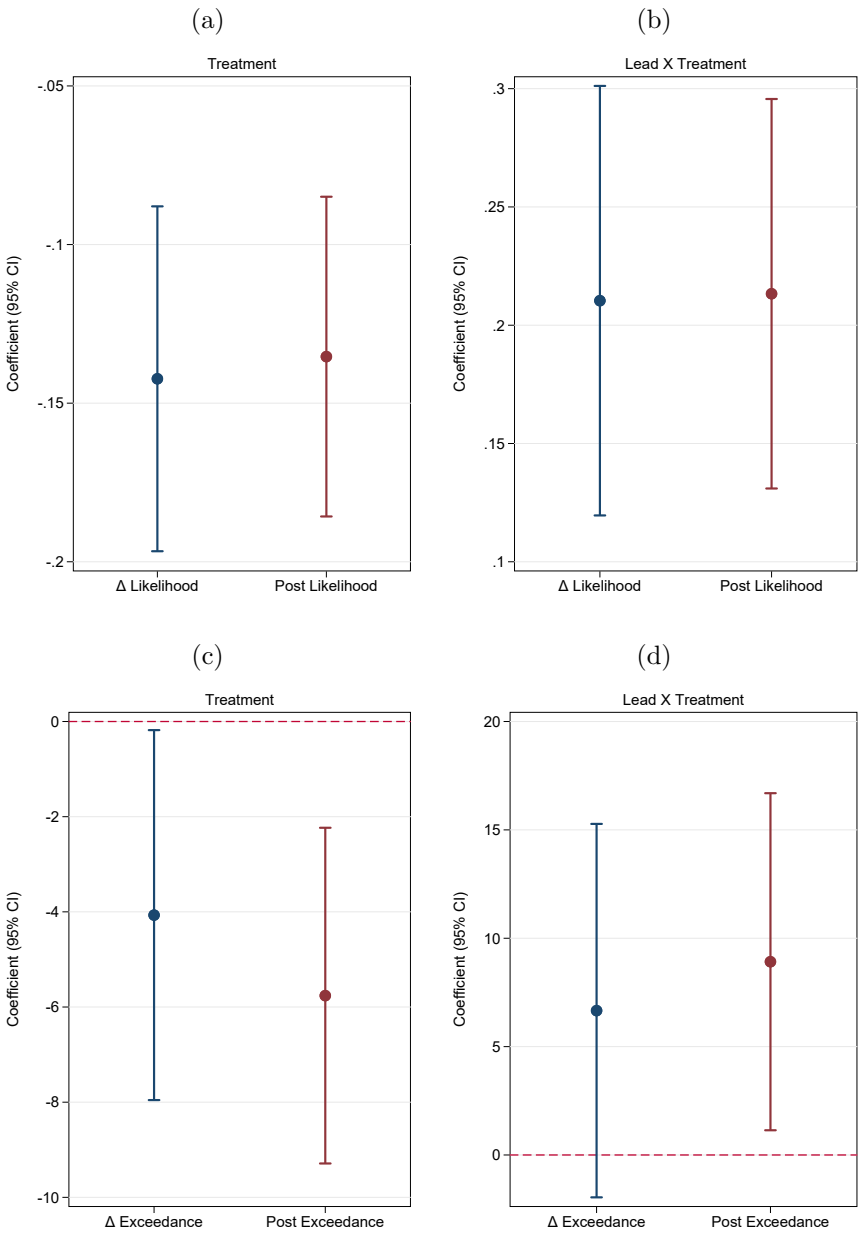


Figure 6: Coefficient point estimates and 95% confidence interval of regressions in Table 7 (estimation results from Equation 1) and Table 10 (estimation results from Equation 2) for outcome variables related to perceived likelihood of LSL presence and perceived number of tap-water samples with lead concentration exceeding the federal limit of 0.1 mg/L. The navy points and whiskers correspond to results in Table 7, while the maroon ones correspond to the results in Table 10. Panels (a) and (b) show results for perceived LSL likelihood, while Panels (c) and (d) show results for perceived tap water lead exceedances. Panel (a) and (c) show estimations for Treatment (i.e., whether the respondent learn about the building-specific service line material) while Panel (b) and (c) show estimations for Lead X Treatment.

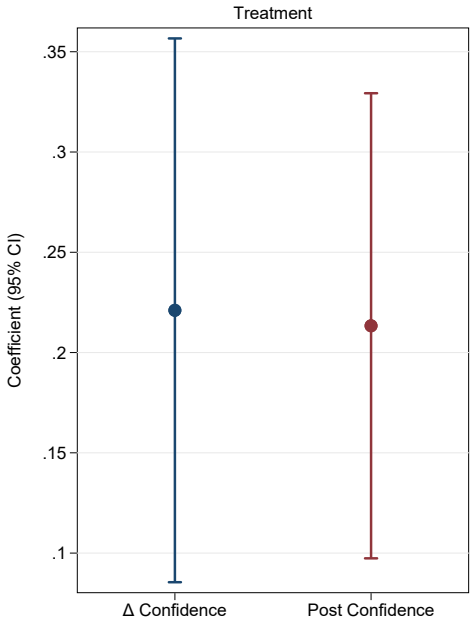


Figure 7: Coefficient point estimates and 95% confidence interval of regressions in Table 7 and Table 10 for confidence about one's perceived water lead exposure. The navy points and whiskers correspond to the result in Column (3) of Table 7, while the maroon ones correspond to the result in Column (3) of Tabel 10.

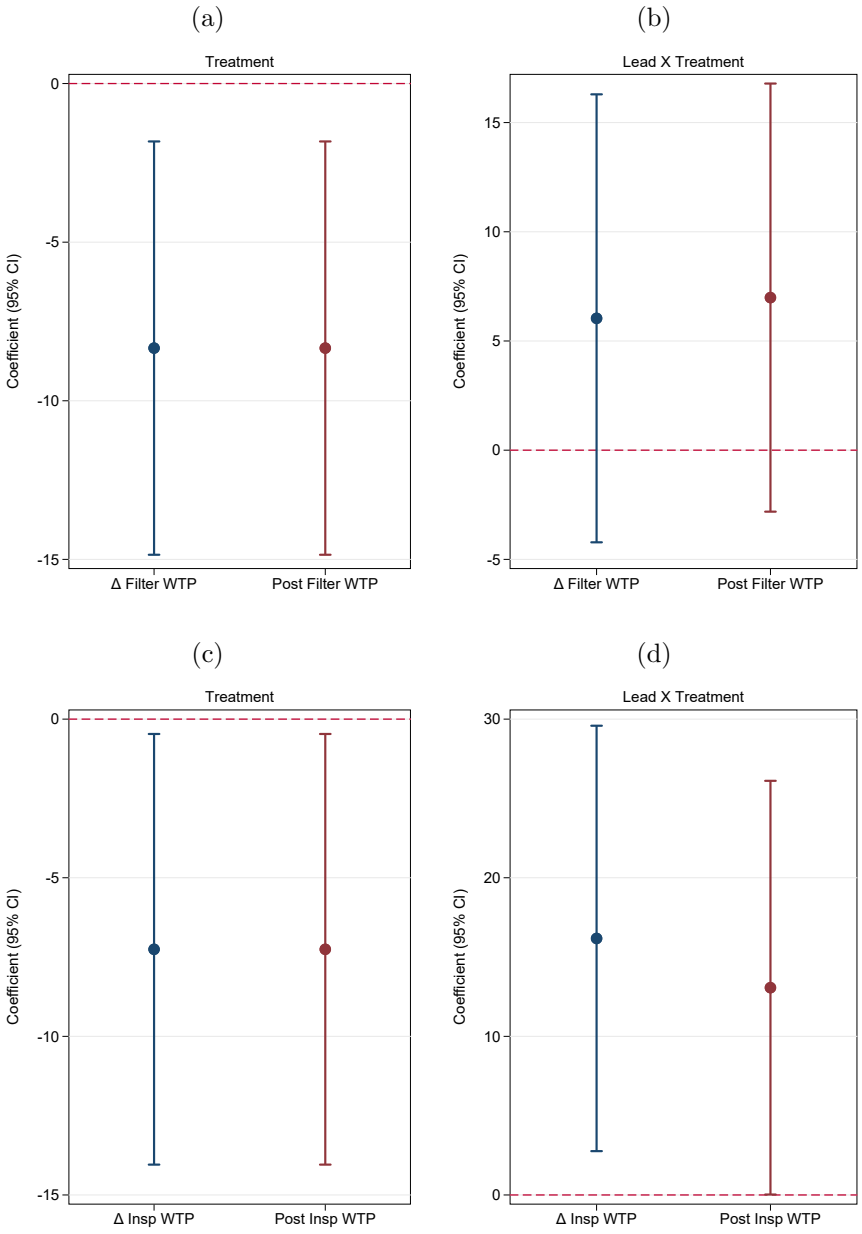


Figure 8: Coefficient point estimates and 95% confidence interval of regressions in Table 7 (estimation results from Equation 1) and Table 10 (estimation results from Equation 2) for outcome variables related WTPs for a lead-filter pitcher and professions tap water lead inspection, respectively. The navy points and whiskers correspond to results in Table 7, while the maroon ones correspond to the results in Table 10. Panel (a) and (b) show results for filter WTP, while Panels (c) and (d) show results for inspection WTP. Panel (a) and (c) show estimations for Treatment (i.e., whether the respondent learn about the building-specific service line material) while Panel (b) and (c) show estimations for Lead X Treatment.

Tables

Table 1: Service line records by city and material

City	Lead	Non-lead	Unknown	Total
Detroit	145 507 39%	85 314 23%	142 409 38%	373 230
Indianapolis	45 875 13%	171 485 50%	129 647 37%	347 007
Milwaukee	70 250 44%	75 756 47%	14 264 9%	160 270
Minneapolis	39 674 33%	79 035 66%	1796 1%	120 505
NYC	128 918 15%	516 840 60%	210 927 25%	856 685
Study Area	385 224 22%	928 430 53%	399 043 25%	1 712 679

Table 2: Respondents by city, material, and treatment assignment

City	Material	Treated	Control	Total
Detroit	Lead	43 31.6%	45 33.1%	88 64.7%
	Non-lead	22 16.2%	26 19.1%	48 35.3%
Indianapolis	Lead	12 16.0%	11 14.7%	23 30.7%
	Non-lead	24 32.0%	28 37.3%	52 69.3%
Milwaukee	Lead	10 17.2%	11 19.0%	21 36.2%
	Non-lead	18 31.0%	19 32.8%	37 63.8%
NYC	Lead	13 4.3%	18 6.0%	31 10.4%
	Non-lead	132 44.1%	136 45.5%	268 89.6%

Table 3: Balance between treatment arms with census benchmarks

Demographics	**Treated n=274**	**Control n=294**	**p-value**	**Census (4 cities)**
Household Size	2.9	2.8	0.787	2.4
Household Age Mean	33.4	32.9	0.700	37.9
Household Income (1–6)	2.8 (*$30–49k*)	2.8 (*$30–49k*)	0.889	3.9† (*$50–74k*)
Respondent Male	54%	50%	0.300	48%
Respondent White	35%	43%	0.057	36%
Respondent Education (1–8)	3.5 (*Some college*)	3.5 (*Some college*)	0.808	3.6† (*Some college*)
Household Health (1–5)	3.6 (*Good*)	3.6 (*Good*)	0.682	NA
Household Chronic Condition	38%	32%	0.137	NA

Notes: Treated/Control are sample means (shares for 0/1). The p-values are from Welch two-sided tests. The last column reports official population-weighted figures across New York City, Detroit, Indianapolis, and Milwaukee. † denotes values imputed from ACS headline shares as documented.

Table 4: Summary of demographics with census benchmarks

Demographics	**Lead n=163**	**Non-lead n=405**	**Census (4 cities)**
Household Size	2.9	2.8	2.4
Household Age Mean	33.5	33.0	37.9
Household Income (1–6)	2.6 (*$30–49k*)	2.8 (*$30–49k*)	3.9† (*$50–74k*)
Respondent Male	49%	53%	48%
Respondent White	29%	43%	36%
Respondent Education (1–8)	3.3 (*Some college*)	3.6 (*Some college*)	3.6† (*Some college*)
Household Health (1–5)	3.6 (*Good*)	3.6 (*Good*)	NA
Household Chronic Condition	31%	36%	NA

Notes: Lead/Non-lead are sample means (shares for 0/1). The last column reports official population-weighted figures across New York City, Detroit, Indianapolis, and Milwaukee. † denotes values imputed from ACS headline shares as documented.

Table 5: Pre-treatment balance between treatment arms

Metrics	**Treated n=274**	**Control n=294**	**p-value**
Belief: lead	58%	56%	0.387
Belief: lead sample	23.7	27.0	0.181
Confidence: lead sample	3.8	3.9	0.825
WTP: pitcher	70.4	65.8	0.219
WTP: inspection	72.9	67.2	0.129

Notes: Values are means. p-values from Welch two-sided t-tests within each material group.

Table 6: Pre-treatment outcomes between treatment × material arms

Metrics	**Lead**			**Non-lead**		
	Treated n=78	**Control n=85**	**p-value**	**Treated n=196**	**Control n=209**	**p-value**
Belief: lead	64%	61%	0.494	55%	54%	0.530
Belief: lead sample	31.4	29.2	0.651	20.6	26.1	0.053
Confidence: lead sample	3.9	3.8	0.751	3.8	3.9	0.650
WTP: pitcher	74.0	66.0	0.203	69.0	65.8	0.481
WTP: inspection	66.5	69.9	0.542	75.5	66.1	0.052

Notes: Values are means. p-values from Welch two-sided t-tests within each material group.

Table 7: Treatment Effects of LSL Information Disclosure (ITTs)

	Δ(Likelihood of LSL)	Δ(Lead Exceedance)	Δ(Confidence)	Δ(Filter WTP)	Δ(Insp WTP)
	(1)	(2)	(3)	(4)	(5)
Lead × Treatment	0.21***	6.66		6.04	16.2**
	(0.046)	(4.39)		(5.22)	(6.83)
Treatment	-0.14***	-4.07**	0.22***	-8.34**	-7.26**
	(0.028)	(1.98)	(0.069)	(3.32)	(3.46)
Lead status = 1	0.042	4.05		-1.14	1.49
	(0.035)	(2.59)		(3.67)	(3.65)
LHS Mean	-0.09	-1.30	0.19	-1.19	-1.42
LHS SD	0.28	21.38	0.83	31.45	35.62
City-FE	Y	Y	Y	Y	Y
N	568	568	568	568	568
adj. R^2	0.11	0.02	0.01	0.01	0.02

Notes. All columns run Equation (1) except Column (3). Change in the outcome variable is regressed on treatment, lead status, the interaction between lead status and treatment, and city fixed effects. Heteroskedasticity-robust standard error in parenthesis (* $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$).

Table 8: Treatment Effects of LSL Information Disclosure (ITTs) w/ Controls

	Δ(Likelihood of LSL)	Δ(Lead Exceedance)	Δ(Confidence)	Δ(Filter WTP)	Δ(Insp WTP)
	(1)	(2)	(3)	(4)	(5)
Lead × Treatment	0.21***	6.50		5.75	16.0**
	(0.046)	(4.41)		(5.18)	(6.84)
Treatment	-0.14***	-4.05**	0.22***	-8.26**	-7.20**
	(0.028)	(1.96)	(0.069)	(3.32)	(3.45)
Lead status = 1	0.041	4.15		-1.26	1.36
	(0.036)	(2.57)		(3.68)	(3.70)
LHS Mean	-0.09	-1.30	0.19	-1.19	-1.42
LHS SD	0.28	21.38	0.83	31.45	35.62
City-FE	Y	Y	Y	Y	Y
Control	Y	Y	Y	Y	Y
N	568	568	568	568	568
adj. R^2	0.11	0.03	0.01	0.01	0.02

Notes. All columns control for household income and the presence of a household member aged 0-10. Heteroskedasticity-robust standard error in parenthesis (* $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$)

Table 9: Effect of the Order of the WTP Questions

	Prior		Change		Posterior	
	Filter	Insp	Filter	Insp	Filter	Insp
	(1)	(2)	(3)	(4)	(5)	(6)
Ask Insp WTP First	-8.12**	-4.38	-0.65	-3.32	-2.27	-4.45
	(3.79)	(3.85)	(2.71)	(3.02)	(2.70)	(2.89)
Pre Filter WTP					0.80***	
					(0.057)	
Pre Insp WTP						0.74***
						(0.058)
LHS Mean	68.05	69.94	-1.19	-1.42	66.86	68.53
LHS SD	44.78	45.33	31.45	35.62	46.88	48.00
City-FE	Y	Y	Y	Y	Y	Y
N	568	568	568	568	568	568
adj. R^2	0.02	0.00	-0.00	0.01	0.58	0.51

Notes. This table shows results from regressing various WTP variables on whether the WTP question for professional tap water lead inspection shows up before that for the lead-filter pitcher in the survey flow for the specific respondent. This order is assigned randomly within each stratum (i.e., city). Heteroskedasticity-robust standard error in parenthesis (* $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$).

Table 10: Treatment Effects of Disclosure: Regressing Post on Prior

	Posterior LSL Likelihood	Posterior Lead Exceedance	Posterior Confidence	Posterior Filter WTP	Posterior Insp WTP
	(1)	(2)	(3)	(4)	(5)
Lead × Treatment	0.21***	8.92**		6.99	13.1**
	(0.042)	(3.96)		(4.99)	(6.64)
Treatment	-0.14***	-5.76***	0.21***	-7.71**	-4.95
	(0.026)	(1.80)	(0.059)	(3.22)	(3.27)
Lead status = 1	0.060*	3.32		-1.65	1.99
	(0.032)	(2.55)		(3.46)	(3.56)
Pre Likelihood	0.60***				
	(0.043)				
Pre Exceedance		0.68***			
		(0.045)			
Pre Confidence			0.56***		
			(0.040)		
Pre Filter WTP				0.80***	
				(0.054)	
Pre Insp WTP					0.75***
					(0.056)
LHS Mean	0.48	24.11	4.03	66.86	68.53
LHS SD	0.31	28.52	0.90	46.88	48.00
City-FE	Y	Y	Y	Y	Y
N	568	568	568	568	568
adj. R^2	0.39	0.56	0.39	0.59	0.51

Notes. All columns run Equation (2) except Column (3). The posterior value of the outcome variable is regressed on treatment, lead status, the interaction between treatment and lead status, the prior value of the same outcome variable, and city fixed effects. Heteroskedasticity-robust standard error in parentheses (* $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$)

Table 11: Treatment Effects of Disclosure: Learning Rate

	Δ(Likelihood of LSL)		Δ(Lead Exceedance)		Δ(Confidence)	
	(1)	(2)	(3)	(4)	(5)	(6)
Lead × Treatment	0.21***		6.66			
	(0.046)		(4.39)			
Perception gap × Treatment		0.20***		8.84**		
		(0.047)		(3.82)		
Abs(Perc Gap) x Treatment						-0.079
						(0.25)
Treatment	-0.14***	-0.021	-4.07**	0.45	0.22***	0.26*
	(0.028)	(0.023)	(1.98)	(2.18)	(0.069)	(0.14)
Lead status = 1	0.042		4.05			
	(0.035)		(2.59)			
Perception gap		0.13***		1.77		
		(0.033)		(2.61)		
Abs(Perception Gap)						-0.021
						(0.20)
LHS Mean	-0.09	-0.09	-1.30	-1.30	0.19	0.19
LHS SD	0.28	0.28	21.38	21.38	0.83	0.83
City-FE	Y	Y	Y	Y	Y	Y
N	568	568	568	568	568	568
adj. R^2	0.11	0.21	0.02	0.03	0.01	0.01

Notes. Columns (1), (3), (5) run Equation (1). Columns (2), (4), (6) run Equation (3). Column (6) uses the absolute value of G_i in Equation (3). Heteroskedasticity-robust standard error in parenthesis (* $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$).

Appendices

A Additional Results

A.1 Additional Figures

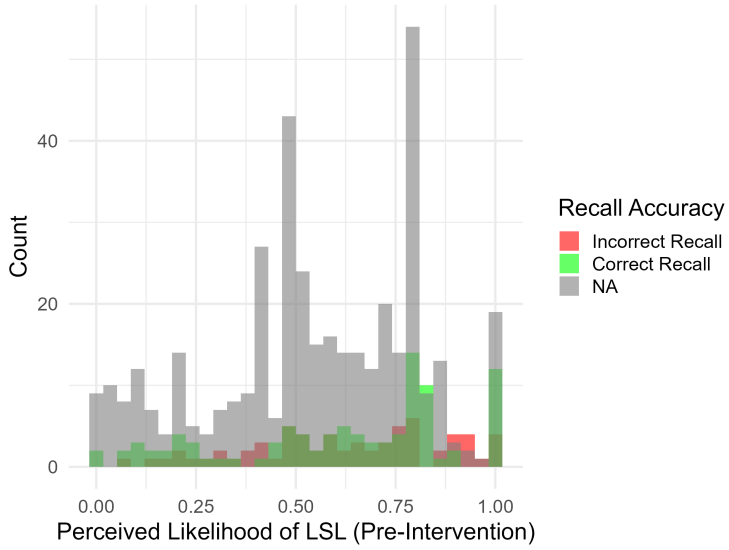


Figure A1: Pre-intervention perceived likelihood of LSL presence by the status of whether respondent could recall correctly about the service line material type information if they self-reported that they had been notified by the utility. Gray represents respondents who either reported no notification or were unable to recall the service line material. Green indicates those who reported being notified and correctly recalled the material type. Red shows the distribution for those who reported being notified but recalled the material type incorrectly.

A.2 Additional Tables

Table A1: Self-reported utility notification and recall by city

City (Sample Size)	Notified	Recall	Correct Recall
Detroit (136)	53	51	26
	39%	38%	19%
Indianapolis (75)	11	10	9
	15%	13%	12%
Milwaukee (58)	24	24	14
	41%	41%	24%
NYC (299)	85	80	49
	28%	27%	16%
Study Area	173	165	98
	30%	29%	17%

Notes. Percentages are calculated based on the number of responses within each city, except for the last row where they are calculated by dividing the total number of responses.

Table A2: Response Screening

Sample Size	Control	Treated	Total	p-value of Chi-square Test for Treatment and LSL Status
All responses with treatment assigned	347	348	695	0.9142
All completes (i.e., attrition removed)	332	330	662	1.0000
Removed if duration < 10 min	306	295	601	0.8155
Removed if incorrect answer to either of the attention checks	294	274	568	0.9807

Notes. (1) The treatment assignment happened right after finishing all pre-intervention questions and before the intervention. (2) The last column reports p-values from Pearson's Chi-square tests used to assess the independence of treatment assignment and LSL status.

Table A3: Other Pre- or Post- Intervention Characteristics

Metric	Treated Mean	Control Mean	Treated N	Control N	p-value
Use tap water for drinking (%)	70.44	68.71	274	294	0.655
Use tap water for cooking (%)	97.08	95.92	274	294	0.451
Any treatment of drinking tap water (%)	73.06	75.25	193	202	0.620
Any treatment of cooking tap water (%)	55.64	51.06	266	282	0.284
Treatment certified to remove lead (%)	71.78	73.08	163	182	0.789
Perceived safety of treated drinking water (1-5)	3.97	4.09	150	165	0.321
Perceived safety of treated cooking water (1-5)	3.93	4.11	192	204	0.089
ADHD diagnosis of >3.5 mcg/dL BLL	3.91	3.90	274	294	0.949
Stroke death in 10,000 adults of >3.5 mcg/dL BLL	2.22	2.23	274	294	0.884
Trust in city's database (0–10)	6.80	6.00	274	294	0.000
Trust in city's database for control vs trust in info received for treated (0–10)	6.88	6.00	274	294	0.000
Perceived effectiveness of water filter (0-10)	7.37	6.90	107	121	0.103

Notes. Water usage questions were before intervention while other questions in the table were **after** all the post-intervention outcome measures, i.e., at the end of the survey.

Table A4: Perceived water safety (on 1-5 scale) regarding lead and E. Coli

	Lead				E. Coli			
	Change		Posteior		Change		Posteior	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lead × Treatment	-0.60***		-0.62***		-0.63***		-0.61***	
	(0.18)		(0.17)		(0.18)		(0.17)	
Perception gap × Treatment		-0.56***		-0.54***		-0.52***		-0.48***
		(0.18)		(0.17)		(0.16)		(0.15)
Treatment	0.32***	-0.0083	0.35***	0.021	0.29***	-0.040	0.28***	-0.032
	(0.090)	(0.095)	(0.087)	(0.090)	(0.087)	(0.089)	(0.083)	(0.085)
Lead status = 1	0.14		0.15		0.070		0.11	
	(0.12)		(0.11)		(0.12)		(0.12)	
Perception gap		0.12		0.22**		0.039		0.15
		(0.11)		(0.11)		(0.100)		(0.098)
Pre lead safety			0.74***	0.75***				
			(0.039)	(0.040)				
Pre E. Coli safety							0.79***	0.79***
							(0.031)	(0.031)
LHS Mean	0.03	0.03	3.63	3.63	-0.05	-0.05	3.66	3.66
LHS SD	0.94	0.94	1.22	1.22	0.92	0.92	1.37	1.37
City-FE	Y	Y	Y	Y	Y	Y	Y	Y
N	568	568	568	568	568	568	568	568
adj. R^2	0.03	0.03	0.49	0.48	0.03	0.03	0.61	0.60

Notes. The questions were asked before and after the intervention, for both lead and E. coli. The response was coded on a 1-5 scale: (1) Very unhealthy for all people, requiring immediate intervention; (2) Unhealthy for all people; (3) Unhealthy for sensitive people such as children, the elderly, or individuals with health conditions; (4) Acceptable, but some health concerns; and (5) No health risk. Heteroskedasticity-robust standard error in parenthesis (* $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$)

B Additional Proofs

B.1 Proof for Proposition 1

Recall

$$U_F(z,\pi) = \frac{1}{\bar{r}} \int_0^z \mathbb{E}_\pi[u(y-r, eW)]\, dr + \left(1 - \frac{z}{\bar{r}}\right) \mathbb{E}_\pi[u(y,W)]\,.$$

Differentiating with respect to z gives

$$\frac{\partial U_F(z,\pi)}{\partial z} = \frac{1}{\bar{r}} \mathbb{E}_\pi[u(y-z, eW)] - \frac{1}{\bar{r}} \mathbb{E}_\pi[u(y,W)]\,.$$

Hence any interior optimum $z^*(\pi)$ must satisfy

$$\mathbb{E}_\pi\big[u\big(y - z^*(\pi), eW\big)\big] = \mathbb{E}_\pi[u(y,W)]\,,$$

which is Equation (5).

For the second-order condition,

$$\frac{\partial^2 U_F(z,\pi)}{\partial z^2} = -\frac{1}{\bar{r}} \mathbb{E}_\pi[u_x(y-z, eW)]\,.$$

By Assumption 1, $u_x > 0$, so the second derivative is strictly negative. Therefore any interior critical point is the unique maximizer.

B.2 Proof for Proposition 2

Recall

$$U_I(z,\pi) = \frac{1}{\bar{r}} \int_0^z V_1(r,\pi)\, dr + \left(1 - \frac{z}{\bar{r}}\right) V_0(\pi).$$

Differentiating with respect to z yields

$$\frac{\partial U_I(z,\pi)}{\partial z} = \frac{1}{\bar{r}} V_1(z,\pi) - \frac{1}{\bar{r}} V_0(\pi).$$

Hence any interior optimum $\tilde{z}^*(\pi)$ must satisfy

$$V_1\big(\tilde{z}^*(\pi), \pi\big) = V_0(\pi),$$

which is Equation (7).

To verify the second-order condition, note that

$$V_1(r,\pi) = \mathbb{E}_\pi\left[\max_{a \in \mathcal{A}} u\big(y - c(a) - r, d(W,a)\big)\right].$$

Let $a^*(W,r)$ denote an optimal action after observing W and paying price r. By the envelope theorem,

$$\frac{\partial V_1(r,\pi)}{\partial r} = \mathbb{E}_\pi\big[-u_x\big(y - c(a^*(W,r)) - r, d(W, a^*(W,r))\big)\big]\,.$$

By Assumption 1, $u_x > 0$, so $\partial V_1(r,\pi)/\partial r < 0$. Therefore

$$\frac{\partial^2 U_I(z,\pi)}{\partial z^2} = \frac{1}{\bar{r}} \frac{\partial V_1(z,\pi)}{\partial z} < 0.$$

Hence any interior critical point is the unique maximizer.

B.3 Proof for Lemma 1

Under Assumption 2,

$$u(x, d) = x + v(d).$$

Therefore

$$V_1(r, \pi) = \mathbb{E}_\pi \left[\max_{a \in \mathcal{A}} \{ y - c(a) - r + v\big(d(W, a)\big) \} \right] = y - r + \widetilde{V}_1(\pi),$$

where

$$\widetilde{V}_1(\pi) \equiv \mathbb{E}_\pi \left[\max_{a \in \mathcal{A}} \{ -c(a) + v\big(d(W, a)\big) \} \right].$$

Likewise,

$$V_0(\pi) = \max_{a \in \mathcal{A}} \mathbb{E}_\pi \big[y - c(a) + v\big(d(W, a)\big) \big] = y + \widetilde{V}_0(\pi).$$

Substituting these expressions into Equation (7) gives

$$y - r + \widetilde{V}_1(\pi) = y + \widetilde{V}_0(\pi) \implies -r + \widetilde{V}_1(\pi) = \widetilde{V}_0(\pi).$$

Evaluating at the optimal bid $r = \tilde{z}^*(\pi)$ yields

$$\tilde{z}^*(\pi) = \widetilde{V}_1(\pi) - \widetilde{V}_0(\pi),$$

which proves Equation (9).

B.4 WTP for water filter with known tap water lead level

Suppose the household knows the realized contamination level w. Then expected utility from bidding z for the filter is

$$U_F(z, w) = \frac{1}{\bar{r}} \int_0^z u(y - r, ew)\, dr + \left(1 - \frac{z}{\bar{r}}\right) u(y, w).$$

Differentiating with respect to z gives

$$\frac{\partial U_F(z, w)}{\partial z} = \frac{1}{\bar{r}} \left[u(y - z, ew) - u(y, w) \right].$$

Hence any interior optimum $\ddot{z}^*(w)$ satisfies

$$u\big(y - \ddot{z}^*(w), ew\big) = u(y, w).$$

Moreover,

$$\frac{\partial^2 U_F(z, w)}{\partial z^2} = -\frac{1}{\bar{r}} u_x(y - z, ew) < 0$$

by Assumption 1, so the interior solution is unique.

B.5 Proof for Proposition 3

Under Equation (8), filter WTP under any belief distribution π can be written as

$$z^*(\pi) = \mathbb{E}_\pi[m(W)], \qquad m(w) \equiv v(ew) - v(w).$$

Before disclosure, beliefs are

$$\pi^0 = qF_1 + (1-q)F_0.$$

Therefore

$$z_F^0 = \mathbb{E}_{\pi^0}[m(W)] = q\mathbb{E}[m(W) \mid S=1] + (1-q)\mathbb{E}[m(W) \mid S=0] = qM_1 + (1-q)M_0.$$

After disclosure of $S = s$, beliefs become $\pi^s = F_s$, so

$$z_F^1(s) = \mathbb{E}_{F_s}[m(W)] = M_s.$$

Subtracting yields

$$\Delta_F(1) = z_F^1(1) - z_F^0 = (1-q)(M_1 - M_0),$$
$$\Delta_F(0) = z_F^1(0) - z_F^0 = -q(M_1 - M_0).$$

This proves Proposition 3.

B.6 Proof for Corollary 1

If $v(d) = -\beta d$ for some $\beta > 0$, then

$$m(w) = v(ew) - v(w) = -\beta ew + \beta w = \beta(1-e)w.$$

Hence

$$M_s = \mathbb{E}[m(W) \mid S=s] = \beta(1-e)\mathbb{E}[W \mid S=s] = \beta(1-e)\mu_s.$$

Substituting into Proposition 3 gives

$$z_F^0 = \beta(1-e)\big[q\mu_1 + (1-q)\mu_0\big],$$

$$\Delta_F(1) = \beta(1-e)(1-q)(\mu_1 - \mu_0), \qquad \Delta_F(0) = -\beta(1-e)q(\mu_1 - \mu_0),$$

as claimed.

B.7 Proof for Corollary 2

If $v(d) = -\beta d^2$ for some $\beta > 0$, then

$$m(w) = v(ew) - v(w) = -\beta e^2 w^2 + \beta w^2 = \beta(1-e^2)w^2.$$

Hence

$$M_s = \mathbb{E}[m(W) \mid S=s] = \beta(1-e^2)\mathbb{E}[W^2 \mid S=s].$$

Using

$$\mathbb{E}[W^2 \mid S=s] = \mu_s^2 + \tau_s^2,$$

we obtain

$$M_s = \beta(1-e^2)(\mu_s^2 + \tau_s^2).$$

Substituting into Proposition 3 gives

$$z_F^0 = \beta(1-e^2)\Big[q(\mu_1^2+\tau_1^2)+(1-q)(\mu_0^2+\tau_0^2)\Big],$$

$$\Delta_F(1) = \beta(1-e^2)(1-q)\Big[(\mu_1^2+\tau_1^2)-(\mu_0^2+\tau_0^2)\Big],$$

$$\Delta_F(0) = -\beta(1-e^2)q\Big[(\mu_1^2+\tau_1^2)-(\mu_0^2+\tau_0^2)\Big],$$

as claimed.

B.8 Proof for Proposition 4

By Lemma 1, the optimal inspection bid under any belief distribution π is

$$\tilde{z}^*(\pi) = \widetilde{V}_1(\pi) - \widetilde{V}_0(\pi).$$

Evaluating this expression at the pre-disclosure mixture $\pi^0 = qF_1 + (1-q)F_0$ gives

$$z_I^0 = \tilde{z}^*(\pi^0).$$

Evaluating it at the post-disclosure distribution F_s gives

$$z_I^1(s) = \tilde{z}^*(F_s).$$

Therefore,

$$\Delta_I(s) = z_I^1(s) - z_I^0 = \tilde{z}^*(F_s) - \tilde{z}^*(\pi^0).$$

This proves Proposition 4.

B.9 Proof for Corollary 3

Under Assumption 3 and linear damages $v(d) = -\beta d$, utility conditional on realized contamination w can be written as

$$y - c(a) - \beta d(w,a) = y - ca - \beta w(1-a) = y - \beta w + a(\beta w - c).$$

Hence the ex post optimal action is

$$a^*(w) = \begin{cases} 0, & w < c/\beta, \\ 1, & w > c/\beta. \end{cases}$$

(At the knife-edge value $w = c/\beta$, either action is optimal.)

If $u_s \le c/\beta$, then $a^*(w) = 0$ for every w in the support of F_s. If $l_s \ge c/\beta$, then $a^*(w) = 1$ for every w in the support of F_s. In either case, observing w does not change the optimal action, so

$$\widetilde{V}_1(F_s) = \widetilde{V}_0(F_s)$$

and therefore

$$\tilde{z}^*(F_s) = 0.$$

If instead $l_s < c/\beta < u_s$, then the support of F_s contains realizations on both sides of the threshold. Because F_s is uniform on $[l_s, u_s]$, both events

$$\{W < c/\beta\} \quad \text{and} \quad \{W > c/\beta\}$$

occur with strictly positive probability. Any fixed action $a \in \{0, 1\}$ is therefore strictly suboptimal on a set of positive probability, whereas an inspected household can condition on w and choose the ex post optimal action state by state. It follows that

$$\widetilde{V}_1(F_s) > \widetilde{V}_0(F_s),$$

so

$$\tilde{z}^*(F_s) > 0.$$

This proves Corollary 3.