

UC Berkeley

CEGA Working Papers

Title

Beliefs, Signal Quality, and Information Sources: Experimental Evidence on Air Quality in Pakistan

Permalink

<https://escholarship.org/uc/item/3z476268>

Authors

Intiaz, Isra

Nakamura, Shotaro

Nasim, Sanval

et al.

Publication Date

2025-01-29

DOI

10.26085/C3FG6D

Series Name: WPS

Paper No.: 247

Issue Date: 1/15/25

Beliefs, Signal Quality, and Information Sources: Experimental Evidence on Air Quality in Pakistan

Isra Imtiaz, Shotaro Nakamura, Sanval Nasim, and Arman Rezaee



CEGA

Center for Effective Global Action

Working Paper Series

Center for Effective Global Action

University of California



This paper is posted at the eScholarship Repository, University of California. http://escholarship.org/uc/cega_wps Copyright © 2023 by the author(s).

The CEGA Working Paper Series showcases ongoing and completed research by faculty affiliates of the Center. CEGA Working Papers employ rigorous evaluation techniques to measure the impact of large-scale social and economic development programs, and are intended to encourage discussion and feedback from the global development community.

Recommended Citation:

Imtiaz, Isra; Nakamura, Shotaro; Nasim, Sanval; Rezaee, Arman (2024): Beliefs, Signal Quality, and Information Sources: Experimental Evidence on Air Quality in Pakistan. CEGA Working Paper Series No. WPS-247. Center for Effective Global Action. University of California, Berkeley.

Beliefs, Signal Quality, and Information Sources: Experimental Evidence on Air Quality in Pakistan*

Isra Imtiaz Shotaro Nakamura Sanval Nasim
Arman Rezaee

December 3, 2024

Abstract

We study how information sources as a signal of service quality shape consumers' beliefs about and demand for air quality forecast services. We provide day-ahead SMS forecasts in Lahore, Pakistan, randomizing whether the forecast is attributed to the government or an NGO. Respondents do not have differential demand by the assigned source but believe the government's forecasts are worse than the NGO's. The results demonstrate that consumers expect lower accuracy from the government, have a limited willingness to pay for accuracy, and prefer the assigned source as they learn about its service quality.

JEL: Q53, Q56, D83, C93

Keywords: air quality, beliefs, environmental information, willingness to pay

*Imtiaz: independent researcher. Nakamura: Federal Trade Commission, Washington DC; email: shotaro.n.nakamura@gmail.com. Nasim: Department of Economics, Colby College. Rezaee: Department of Economics, University of California, Davis. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the Federal Trade Commission, or its Commissioners. This study is funded by grants from the International Growth Center (IGC) and faculty funding from Colby College and UC Davis. This study has IRB approval in the US (UC Davis IRB-1336133-1) and Pakistan (LUMS IRB 02242021). This study was pre-registered with AEA (AEARCTR-0011489).

1 Introduction

Information about the environment is a public good, which standard economic theory predicts markets would under-supply (Samuelson 1954). Limited access to air quality information would be costly to citizens in developing economies given significant health and productivity damages from, and seasonal variations and uncertainty about, ambient air pollution (WHO 2021; Barwick et al. 2024). According to Greenstone et al. (2024), the average loss of life years due to high concentrations of particulate matter (PM_{2.5}) in South Asia is 3.5 years, as opposed to months in the U.S. and Europe.¹ Not only is air quality in places like Lahore, where we conduct our study, an order of magnitude worse (PM_{2.5} concentration of $124.8\mu\text{g}/\text{m}^3$ on average in 2022) than the U.S. ($7.8\mu\text{g}/\text{m}^3$), but it is also driven by seasonal fluctuations based on factors like residual crop burning (Lan et al. 2022).²

Since markets under-supply public goods, governments often assume the responsibility of providing their desired levels. Yet, some government agencies fall short as they face resource and capacity constraints or perverse incentives to obscure the true extent of environmental degradation (e.g., Ghanem and Zhang 2014; Mu et al. 2024). As a result, there is often a large and unmet demand for air quality information and effective mitigation measures against air pollution, particularly in developing economies (e.g., Ahmad et al. 2022; Freeman et al. 2019; Ito and Zhang 2020). One solution is to have competing non-governmental services. Various stakeholders, such as citizen-led advocacy groups and international and bilateral agencies, have begun providing air quality information in developing economies, leading to increased awareness and improved local air quality (Jha and La Nauze 2022).

However, it is unclear how citizens respond to multiple information sources and how their responses ultimately affect the accessibility and quality of the

¹The measure is defined as the average life years gained by a reduction in the concentration of PM_{2.5} to the 2021 WHO standards ($5\mu\text{g}/\text{m}^3$).

²Average concentration measures are authors' own calculations based on U.S. EPA sources: <https://www.airnow.gov/international/us-embassies-and-consulates> <https://www.epa.gov/air-trends/particulate-matter-pm25-trends>.

public good (Bergstrom et al. 1986; Coury et al. 2024). While citizens may have more points of contact with different providers, they may not consider different sources as perfect substitutes if they have differential beliefs and preferences by the source. A body of work addresses whether public-private competition helps improve public service quality in developing economies (Andrabi et al. 2017; Das et al. 2016; Muralidharan and Sundararaman 2015). Little is known, however, about how consumers’ beliefs and preferences for a service provider evolve when there is competition to public provision. Previous papers, such as Acemoglu et al. (2020), Khan et al. (2021), and Dhinakar Bala et al. (2024), study how information helps build trust in the state and its capacity. Yet, they have not addressed how a provider, holding constant service quality, shapes beliefs about and preferences for public goods in developing economies.

In this article, we study how consumers form beliefs about air quality information as they infer the service quality of two competing sources: the government and an NGO run by concerned citizens. We address the following research questions. First, is there an unmet demand for air quality information services in developing cities, particularly among lower-middle-income citizens? Second, does attribution to a particular source affect the demand for air quality information? Third, what mechanisms explain the differential demand, or the lack thereof, by the source? In particular, how does attribution affect the relative preferences between sources and beliefs about service quality? Fourth, what do our results imply about the costs and benefits of competing sources in terms of access to information and welfare?

We address these questions via a randomized controlled trial (RCT) across 1,010 households in working-class neighborhoods of Lahore, Pakistan. We identify that government services on air quality information are limited despite severe ambient air pollution, which has prompted several alternative providers to publish independently collected air quality readings. We reflect this information environment to our intervention by developing an ensemble forecast model of day-ahead air quality levels using data from multiple sources.

In our intervention, we send identical daily air quality forecasts to the sampled residents via a mobile Short Message Service (SMS) for free and randomly vary the source to which we attribute the forecasts. The intervention is unique from the set of existing information services in that it a) provides day-ahead forecasts and b) makes it easier to acquire information for average citizens, who report to have limited access at baseline.

We randomly vary the source to which we attribute the information source in our experimental sample. In one arm, respondents are told that the daily forecasts are constructed using data from the Punjab Environmental Protection Department (EPD), a government agency responsible for reporting on air quality. In the other arm, respondents are told that the forecasts are constructed using data from an NGO called Pakistan Air Quality Initiative (PAQI), a leading local organization that is still relatively unknown to the general population. We do not have a pure control arm in which we do not provide air quality forecasts, as the effect of information provision has been extensively studied by Ahmad et al. (2022) in a similar context. Instead, we focus on the effects of attribution to one source relative to another.

The experimental setup allows us to measure whether the forecast recipients value air quality information and how they perceive and trust the sources from which the information comes. We conduct incentivized games in which we measure their willingness to pay for air quality forecasts, elicit their beliefs about air quality levels and the accuracy of the SMS forecasts, and observe changes in their preferences between sources. We map our empirical hypotheses to a conceptual framework where we define how a consumer's demand for information depends on the source. We specify that utility depends not only on consumers' beliefs about the state of air quality but also on their beliefs about the accuracy of information and other dimensions of service quality.

We find that residents of working-class neighborhoods in Lahore have high demand for air quality forecasts, yet not differentially so based on the salient information source. The average willingness to pay for two additional months of service after experiencing it for free is 238 Pakistani Rupees (PKR) as

measured through the Becker-DeGroot-Marshak (BDM) method, equivalent to a month of standard mobile and data services. The high willingness to pay suggests robust demand for air quality information and satisfaction with the service we provide, which we confirm via stated preference measures in the endline survey. Yet, we do not find statistically significant differences in either their willingness to pay.

We hypothesize that sources affect beliefs and preferences on various attributes that constitute consumers' utility from the SMS forecasts. As such, we test the treatment effects on the following attributes: a) consumers' beliefs about the state of air quality, b) their beliefs about the accuracy of information, and c) their relative preferences between sources that may include factors beyond the accuracy of information. We systematically provide evidence on these three channels as follows.

First, we do not find evidence that the sources differentially affect beliefs about air quality levels, as measured through the recipients' errors in incentivized forecast games. We also do not find that avoidance behaviors and policy preferences are different between treatment arms. These results align with the fact that actual forecast values are identical between treatments by design, suggesting that the source alone does not lead to differential changes in beliefs about air quality levels and avoidance behavior. We also rule out an alternative explanation that the information does not contain any useful signals and does not affect recipients' avoidance behavior; using time-use data, we find correlational evidence that recipients respond to air quality information by spending less time outdoors on more polluted days.

Second, we find that information associated with the government is believed to be less accurate, despite the non-differential willingness to pay. Through a set of incentivized forecast games, we find that recipients in the government arm expect a 12% higher error in the SMS forecasts than those in the NGO arm despite receiving identical forecast values. This result is consistent with the idea that recipients do not value forecasts' precision on the margin; instead, they value other attributes more, such as ease of access and reliability.

Third, we find that the recipients' preferences between sources shift significantly as a result of random assignment to a source. We measure relative preference for sources via a donation game in which forecast recipients split a fixed sum for donation between the government and NGO. At baseline, most recipients choose to split their endowments equally between the government and NGO. Yet, at the endline, recipients in the government arm favor the government over the NGO by 75:25, and vice versa for those in the NGO arm. Stated preference measures on accuracy, reliability, and overall approval improve in favor of the assigned source. Our findings suggest that preferences for a source are relatively malleable in a frictional market for information services.

We also estimate heterogeneous treatment effects along pre-specified dimensions of recipients' prior beliefs collected at baseline. We do not find statistically significant heterogeneous effects on the demand for air quality information based on recipients' baseline beliefs. However, we also find heterogeneous effects on the endline air quality forecast errors by their baseline errors; for recipients in the government arm relative to the NGO one, having a 100% larger baseline forecast error is associated with having 26% higher endline forecast errors. These results imply that, even when attributions to information sources do not meaningfully affect the demand for air quality information, consumers with less accurate beliefs about air quality update their beliefs more slowly when they are assigned to the government source, which they believe to have lower quality.

Our work makes several contributions to the economics literature, addresses concerns over external validity, and provides policy insights. In particular, we focus on the following three themes: the demand for environmental goods and services, public-private competition over public service provision, and belief formation under competing signals.

First, our paper contributes to the body of work on the determinants of demand for environmental goods and services in developing economies. Existing studies in China and South Asia largely show unmet demand for avoidance strategies against air pollution (Ahmad et al. 2022; Barwick et al. 2024; Free-

man et al. 2019; Ito and Zhang 2020). Where there is limited demand for avoidance strategies such as in Greenstone et al. (2021), an emerging body of work suggests the importance of (mis)beliefs in shaping consumers' demand (Chowdhury et al. 2024). Our work is also closely related to studies from developed economies on the demand for more accessible forms of information services, such as SMS and mobile apps, than government monitors focused on regulatory purposes. Coury et al. (2024) and Zivin et al. (2024) show that private demand for low-cost monitors in the U.S. is concentrated in wealthier and less polluted areas, exacerbating the inequity in access to environmental public goods. Taken together, our work highlights the importance of consumer beliefs when addressing market failures in environmental public service provision, particularly in disadvantaged communities.

Second, we provide insights into how consumers respond to public-private competition in public service provision. Our work is motivated by evidence on how government agencies may face perverse incentives to under-report air pollution (Ghanem and Zhang 2014). Our work is also related to studies on the effects of public-private competition on public service quality (Andrabi et al. 2017; Das et al. 2016; Jha and La Nauze 2022; Muralidharan and Sundararaman 2015) and more generally to studies on citizens' trust in the state and its capacity (Acemoglu et al. 2020; Khan et al. 2021; Dhinakar Bala et al. 2024). We build on the literature by focusing on consumer's internal processes by providing detailed belief measures and addressing the role of identity in consumer preferences (Akerlof and Kranton 2000). Our findings suggest that policymakers may be able to induce welfare gains by providing information regardless of the source and by improving the ease of access rather than accuracy, at least on the margin.

Third, our work relates to the literature on beliefs and trust in information sources (Gentzkow et al. 2023; Baysan 2022; Chopra et al. 2022). We shed light on a) the role of beliefs and trust in shaping the demand for environmental information and b) the importance of prior beliefs and conditions under which beliefs about the state of the world and preferences for information services

may diverge. We find that consumers prefer the experimentally assigned source over the alternative, suggesting that their beliefs are relatively malleable to new signals in the short run. However, one may be concerned that preferences over information sources get “locked in” either through platform design or polarization in beliefs (Bowen et al. 2023; Shapiro and Varian 1999). Further work is needed to study effects of competing information sources on consumers’ beliefs in the long run.

The rest of this article is organized as follows: In section 2, we discuss the landscape for air quality information in Lahore. In Section 3, we provide details on the experimental design. We define the pre-specified outcome variables in Section 4 and the identification strategy in Section 5. In Section 6, we present our empirical results. We then introduce a conceptual framework and use it to map our empirical results and pinpoint potential mechanisms. We conclude with policy implications in Section 8.

2 Context

2.1 Air quality information sources

Punjab Province’s Environmental Protection Department (EPD) is the primary regulatory body in Lahore with a mandate to protect the environment and meet the national environmental quality standard, including providing environmental information (The Parliament of Pakistan 1997). EPD is expected to publish daily reports of scheduled pollutants.

Daily readings by EPD, however, are hard to access and often unavailable for reasons that are unclear. EPD only makes public the daily readings in English on their website and does not, to our knowledge, publicize their data in a consistent and timely manner in an accessible form such as social media.³ Those interested in the readings would have to access the website and download a daily PDF, an example of which is shown in Appendix Figure A.1. This is

³The daily reports are posted at <https://epd.punjab.gov.pk/aqi>

a significant hurdle for an average citizen with limited access to the internet and the lack of proficiency in English.⁴ Furthermore, EPD’s PM2.5 readings are missing at a much higher frequency than other sources. Figure 1 shows that the readings were missing for most of December 2022, usually one of the worst air quality periods due to the seasonal smog. As a result, we find that only around 9 percent of the working-class citizens in Lahore report having accessed air quality readings from EPD, as shown in Appendix Table A.1.

Limited information from EPD has led citizens’ initiatives, like the Pakistan Air Quality Initiative (PAQI), to collect and publish their own data to fill the void left by government services. Established in 2016, PAQI crowdsources several low-cost air quality monitors (IQAir and PurpleAir) installed at private homes, businesses, and educational institutions. PAQI, among other operators, uploads their PM2.5 readings to an online platform named AirVisual. The platform reports both monitor-level and city-level readings at the hourly and daily levels, going back as far as one month.⁵ PAQI also has a Twitter account that disseminates daily readings from Lahore.⁶

Yet, the vast majority of the working-class population is unaware of PAQI’s or other sources’ initiatives, as they may be constrained in their smartphone’s data capacity to access specialized apps for air quality or Twitter.⁷ Appendix Table A.1 also shows that approximately 9% of our sampled households stated to have accessed air quality readings from the AirVisual app at baseline. Other sources of air quality information exist for Lahore but are less known than the EPD or PAQI or are equally challenging for the average citizen to access. The most prominent among such sources is called AirNow, a high-quality monitoring system operated by the U.S. Consulate General in Lahore.⁸ The

⁴According to GSMA (2024), 53% of men and 33% of women have access to mobile internet. However, much smaller shares (26% for men and 11% for women) access mobile internet on a daily basis for multiple use cases.

⁵e.g., Lahore and Lahore American School

⁶@LahoreSmog

⁷Anecdotally, Twitter is considered to be an upper-middle-class social network in Pakistan, while Facebook, WhatsApp, and voice-based social media services that require less data are popular among working-class populations.

⁸The U.S. Consulate General in Lahore hosts an air quality monitor funded by the U.S.

Consulate shares their readings on their website and on Twitter, but, to our knowledge, does not actively engage in other forms of dissemination.

2.2 Multiple, conflicting information sources

Information friction and uncertainty about service quality may lead to a world in which air quality readings differ significantly between information sources. We find that the service quality of EPD and PAQI differ significantly in terms of data availability and correlation with an independent U.S. Consulate measure. Appendix Table A.2 shows the summary statistics of daily $PM_{2.5}$ readings by the air quality information source before and during our intervention. Even though EPD values are higher than the PAQI values during the pre-intervention period with high pollution, EPD has 36 fewer observations due to non-reporting. Appendix Tables A.3 and A.4 also show that the EPD readings are less correlated with, and have higher deviations from, the U.S. Consulate measures than the PAQI ones do.

As a result, consumers must decide whether, and from which source, they want to consume air quality information. To do so, they need to gauge the true extent of air pollution and assess the veracity, reliability, and other relevant attributes of a given information source. Inferring service quality, however, may be difficult in a context where there is political polarization and concerns for misinformation in social media (Davies 2023; Hirshleifer et al. 2023). In fact, the Provincial government has signaled its desire to push back by suggesting possible legal action against “fake air quality data” (Raza 2021). Furthermore, EPD reports now contain a disclaimer that “[any] other data from any source presenting ambient air quality of any city of Punjab is neither verified nor approved by the EPA Punjab” as shown in Appendix Figure A.1.

EPA. The AirNow International program places air quality monitors at U.S. embassies and consulates in mainly developing countries and provides hourly historical readings of $PM_{2.5}$ concentration. The monitor is located within the U.S. Consulate’s compound in Shimla Hills, Lahore. The standards for the monitors installed are provided at https://www.epa.gov/system/files/documents/2022-12/List_of_FRM_and_FEM.pdf.

3 Experimental design

To address how attribution to a source affects consumers’ beliefs and demand for information, we conduct an intervention in which we hold the signal quality constant and vary the information sources that we make salient. We randomly sample 1,010 households into two treatment arms: government (EPD) and NGO (PAQI). We optimize power to detect differences between the treatment arms and do not have a pure control arm. We provide the same day-ahead forecast and the day-of readings to the two treatment arms. In each message, we make salient the information source to which recipients are assigned. We capture changes in recipients’ beliefs and preferences about air quality and information sources through lab-in-the-field games. The timing of our intervention and surveys is shown in Appendix Figure A.2.

3.1 Defining the measure of ground truth

The existence of multiple and often conflicting information sources creates a conceptual and empirical challenge: defining the ground truth of air quality levels. Because our research questions revolve around the role of information sources in consumers’ beliefs about some objective measures of true air quality levels, we must choose an independent source that is neither the EPD nor PAQI to be our ground truth. We choose the U.S. Consulate monitor as such a measure because its readings are likely accurate measures of the ground truth at the monitoring station. The U.S. Department of State requires consular monitors and data quality to be of the same technical standard as ones used by the U.S. EPA domestically (White 2018).

The measure for which we construct a forecast model is the daily average concentration of PM_{2.5} (in $\mu g/m^3$) at the U.S. Consulate. We construct a single daily measure and abstract away from spatial variation in air quality, as our respondents reside in a relatively confined, single subsection of the city.⁹

⁹Previous studies have shown that exposure to air pollution vary by socioeconomic status within cities (e.g., Hsiang et al. 2019). However, this spatial variation is not a concern for

The daily based on hourly readings between 12:00 AM and 4:00 PM.¹⁰ The time window is selected so that the research team can collect the day’s readings, estimate the next-day forecast, and send the SMS with the day’s readings and the next-day forecast to our sample households between 6:00–8:00 PM. We identified via our pilot that this timing allows SMS recipients to make plans for the next day based on the SMS forecast.

3.2 Forecast model

We construct an ensemble prediction model of the U.S. consulate readings for the next day ($t + 1$). The model’s objectives are twofold: 1) to provide the most accurate forecast possible given the available information, and b) to ensure that information from both EPD and PAQI are used to construct the forecast. Ensemble modeling allows us to selectively attribute to a source (EPD or PAQI) while holding the actual forecast values constant.

First, we construct four forecast models, all of which predict the ground truth, as inputs for the ensemble forecast model. Each model contains readings from only one monitor as inputs. The readings data on the right-hand side are the $t - 6$ to t lagged readings of either the U.S. Consulate, EPD, PAQI, or the Urban Unit. Since SPRINTARS already provides predictions based on their model, we simply take its $t + 1$ forecast. Each model, except for SPRINTARS, also uses historical meteorological readings and weather forecasts for $t + 1$ as inputs. For each of the input models, we use an adaptive Lasso framework and predict $j + 1$ PM2.5 concentration using a model trained on data from Day 1 to Day j , for j going from Day 20 to t . We have $t - 20$ out-of-sample forecasts, the last of which is for Day $t + 1$, for each model.

our randomization and intervention, as the accuracy of ground truths or forecasts does not depend on treatment arms.

¹⁰We rely on other sources when the U.S. Consulate monitor readings are unavailable. When the U.S. Consulate readings are missing, we use the Urban Unit readings, which are also based on a high-quality monitor (BAM-1020 by MET). If the Urban Unit is also missing, we use PAQI, whose readings are always available. As of 24 May 2023, the U.S. Consulate monitor is missing readings for 16 of the 97 intervention days. Out of 16 days where the U.S. Consulate is missing data, the Urban Unit is missing on 4 days.

We then combine the input forecasts to construct an ensemble model. We estimate the root-mean-square error (RMSE) of each model over the period in which we have forecasts. We then weight the forecast based on the sum of RMSE across five models to their own. The ensemble forecast is the weighted sum of the individual forecasts.

3.3 Sampling

We conduct our intervention in lower-middle-class neighborhoods of National Assembly Constituencies 123 and 124 in northern Lahore. We divide the two constituencies into $200\text{m} \times 200\text{m}$ grids and randomly select 100 blocks, weighted by population density. We then sample 1,010 households from the selected block centroids following the left-hand rule: survey every ten households by spiraling out from the centroid counterclockwise.

3.4 Randomization

The sampled households are divided into two treatment arms. In T1, SMS forecasts are attributed to a government agency (EPD), while in T2, they are attributed to the NGO (PAQI). We do not have a pure control group that does not receive SMS forecasts, as the purpose of this study is to understand the effect of the source, holding constant service quality and forecast values.

We stratify the randomization process into the two treatment groups on a set of baseline variables that a) we considered as potential outcome variables, b) proxies of potential outcome variables that we were unable to collect at baseline due to the experimental design, c) some dimensions of heterogeneity that were considered pre-intervention, and d) the household asset index. We use the optimal-greedy algorithm and generate blocks using the Minimum Volume Ellipsoid (MVE) estimator. We are primarily concerned about balance on outcome variables at baseline and the “take-up” in terms of exposure and comprehension of our SMS forecast messages. We follow the advice from

Athey and Imbens (2017) that each block contains two units per treatment arm. We then split the subjects into T1 and T2.

3.5 Intervention: SMS forecast messaging

The main element of our intervention is the daily provision of the day-ahead (i.e., $t + 1$) forecasts of PM 2.5 measures in $\mu g/m^3$ via SMS. The daily messages are sent from the beginning of the intervention on 18 February, 2023, and continues through the endline survey. The daily messages as part of the intervention ends on 20 June, 2023. All messages were in Urdu. The English translation of the messages on, for instance, 20 May 2023 is as follows:

- T1: "Actual Air Quality (PM 2.5) on 20-05-23: 106
Air Quality Forecast (PM 2.5) for 21-05-23 using data From Punjab Government (EPD): 120
- T2: "Actual Air Quality (PM 2.5) on 20-05-23: 106
Air Quality Forecast (PM 2.5) for 21-05-23 using data From NGO (PAQI): 120

Figure 2 also shows screenshots of the daily messages for T1 and T2. The messages are sent around 6:00–8:00 PM daily, after collecting the day’s data and estimating the forecast for $t+1$. The daily messages also contain the readings from time t . Because the text messages are sent from the same number every day, it is easy to compare the forecast values for Day t provided on Day $t-1$ to the realized value provided on Day t . The subjects also receive an introductory message before the start of the daily SMSs and a reminder message every two weeks over the course of the intervention, as discussed in further detail in Appendix Section B.1.

4 Primary outcome variables

We identify four primary outcomes, constructed using incentivized games, with which we test five primary hypotheses.

4.1 Demand for air quality information as the willingness-to-pay (WTP) for SMS-based air quality forecasts

In the endline survey, we ask for the respondent's willingness to pay for the SMS-based air quality forecast messages for two additional months. The outcome is defined as the amount respondents are willing to pay in PKR. We elicit respondents' willingness to pay for the SMS forecast using the Becker-DeGroot-Marshak (BDM) method (Becker et al. 1964). In the prompt, we make the experimentally assigned source salient by reminding them that the forecast is built using data from the said source. The bid ceiling is PKR 400.

4.2 Beliefs about air quality levels as the absolute error of incentivized $t + 1$ forecast of PM2.5 concentration

In baseline and endline surveys, we ask respondents guess the air pollution level on the next day. We show respondents a table containing the average, minimum, and maximum of the average daily PM2.5 concentration over the last calendar week and ask them to forecast tomorrow's average PM2.5 concentration. Respondents receive PKR 250 if their guess falls within 5% of the actual levels, PKR 150 if within 10%, and PKR 50 if within 20%. The outcome is defined as the absolute difference between the actual PM2.5 concentration and the respondent's forecast, divided by the actual PM2.5 concentration.

4.3 Perceived accuracy of the source as the absolute difference between own and SMS forecasts

In the endline survey, we not only ask respondents to forecast the actual PM2.5 concentration for tomorrow but also the value of our SMS forecast. The guess is financially incentivized, as in the guess for the actual PM2.5 concentration for tomorrow. The outcome is defined as the absolute difference between the respondent’s guess of the PM2.5 forecast generated by our model and their own forecast for $t + 1$.

4.4 Preference for information source as the share of donations to government vs. NGO

In baseline and endline surveys, we offer an opportunity to donate PKR 100 between two sources for environmental protection: a government institution and PAQI. The outcome is defined as the share of PKR 100 donated to a government agency for an environmental cause, as opposed to the NGO.

5 Identification strategy

5.1 Exogenous variable

Our exogenous variable is treatment assignment between the arm where the government (EPD) was made salient as the source, as opposed to the NGO (PAQI). For expository purposes, we refer to being in the government arm as being in the “treatment,” and the NGO arm as the “control.” Let Z denote treatment assignment as a vector, whose inputs are equal to 1 if the respondent is assigned to the government arm and 0 if assigned to the NGO arm.

5.2 Test of positive willingness to pay for air quality information

First, we conduct a t-test to see if the willingness to pay for the SMS forecasts is higher than 0. We pool the two treatment arms.

5.3 Treatment Effects

We estimate the treatment effects between subjects using one of the following equations:

$$Y_i = Z_i' \beta + \gamma Y_{0i} + \mathbf{X}_i' \boldsymbol{\delta} + \varepsilon_i \quad (1)$$

$$Y_i = \alpha + Z_i' \beta + \mathbf{X}_i' \boldsymbol{\gamma} + \varepsilon_i \quad (2)$$

We use Equation 1 for outcomes that have baseline measures and Equation 2 for those that do not. The matrix \mathbf{X} includes control variables selected through a double-post-selection method using LASSO, as in Belloni et al. (2014). Given that we are agnostic as to which information source is more likely to shift beliefs, preferences, and beliefs related to air quality, our hypothesis tests are two-tailed: $\beta \neq 0$.¹¹

6 Results

6.1 Checks on balance

We test for balance of key demographic and baseline outcome variables between the two treatment arms. The statistics we present include means for the two treatment arms, differences between the two treatment arms, and t-tests of

¹¹We also pre-specified a treatment-on-the-treated identification strategy in the pre-analysis plan. However, we do not find significant first-stage results and therefore exclude this identification strategy from our analysis.

the null hypothesis of zero difference. Appendix Table A.1 shows the balance on the variables used in the blocking procedure. We do not find statistically significant differences in any primary outcomes for which we have baseline measures or other variables over which we stratified our randomization.

6.2 Pre-specified outcomes

Table 1 shows the coefficients and their standard errors of the intend-to-treat estimates for the five pre-specified primary hypotheses using post-double-selection LASSO. Here, by “treatment,” we mean being assigned to the government arm instead of the NGO arm. Table 1 also shows the p- and q-values of the corresponding columns. In the following subsections, we focus our analysis on the four pre-specified primary outcomes and five hypotheses. We then complement the findings with non-primary outcomes and analyses.

6.3 Willingness to pay for air quality information

We find that the respondents have a high willingness to pay, but not differentially between treatment arms. Column 1 in Table 1 shows that the respondents are willing to pay PKR 238 for two months of air quality forecast services. This amount is roughly equivalent to a month of popular prepaid mobile and data services, often referred to as the “social” bundle by major carriers in Pakistan. For example, the Social Plus plan by Jazz includes 10Gb of data, 300 minutes of calls in-network, 50 minutes out-of-network, and 1,000 SMS messages and is priced at PKR 260 as of August 2023.¹² Figure 3 also shows the distribution of the willingness-to-pay for air quality forecasts as demand functions, indicating considerable heterogeneity.

However, there are no statistically and economically significant differences between the treatment arms in their willingness to pay for the forecasts. Column 2 in Table 1 shows that those assigned to the Government arm are willing

¹²<https://jazz.com.pk/prepaid/monthly-social-plus>

to pay only PKR 0.33 more on average, and the difference is not statistically significantly different from zero. The small coefficient and standard error exclude any economically meaningful difference between the two treatment arms. We also do not find evidence that the distributions of the willingness to pay are significantly different between treatment arms. Appendix Table A.5 shows that the densities of the 50-rupee bins are not statistically significantly different in 7 out of 8 bins. Appendix Table A.6 also shows that we fail to reject the null in the Kolmogorov-Smirnov and Wilcoxon rank-sum tests.

What would explain the high demand for information, yet no distinction in their differential willingness to pay by the information source? One possibility is that our experiment failed to make the information source salient or that information sources do not matter in forming or shaping recipients' beliefs and preferences. We rule out this possibility in the following sections. Another possibility is that the source affects recipients' beliefs and preferences for air quality information that factor into their utility from consuming the air quality forecast. But such effects may not be picked up in the aggregate willingness-to-pay measure if a) the effects are symmetric between treatment arms or b) the magnitude of the effects is small at the margin.

As such, we follow Equation 9 and highlight three attributes that we hypothesize are key to the recipients' willingness to pay for air quality information and test how the treatment assignment affects them. The first attribute is the recipients' beliefs about the state of air quality; if the treatment affects recipients' beliefs about air quality levels and their ability to forecast it, then that may affect their willingness to pay for future air quality information. The second is the recipients' beliefs about the accuracy of information. The third is the recipients' relative preference between the two sources. We examine these three attributes systematically in the following subsections.

6.4 Beliefs about the state, i.e., air quality levels

We find that different information sources do not lead to differential beliefs about air quality levels nor improve forecast ability in one treatment group relative to another. We measure the respondents' beliefs about air quality via incentivized forecasts. We would expect differences in the forecast error by treatment arm if information sources affect the magnitude by which recipients update their beliefs about the state of air quality toward our SMS forecast. Column 3 on Table 1 shows that those assigned to the government arm have, on average, a five percentage-point higher forecast error than in those in the NGO arm, although the difference is not statistically significant. The magnitude is also small relative to the NGO-arm average error of 73% of the actual reading.

Other measures of beliefs about air pollution levels confirm this finding. First, we do not find statistically significant treatment effects on other definitions of air quality forecast, such as in level and absolute differences, as shown in Appendix Table A.9. Second, we also do not find significant differences between treatment arms in stated measures of concern about air quality. Appendix Table A.10 shows statistically insignificant results on the Likert-scale measure of concern about air quality and the number of days in the last week that the respondents believed to have had good air quality.

Furthermore, the information sources do not differentially affect avoidance behaviors and policy preferences. These results are shown in Appendix Tables A.11 to A.15. We do not generally find statistically significant differences by treatment arms on time spent outdoors, policy preferences for air quality over other issues, and their willingness to file a complaint to the local government. Appendix Table A.12 shows, however, that the recipients in the Government arm are 3.5 percentage points less likely to report to have a mask and 4.3 percentage points less likely to have used a mask in the last week. We refrain from over-interpreting these results due to the concerns about multiple hypothesis testing and the lack of effects on similar outcomes such as time use.

One may be concerned that recipients do not update their beliefs about air

quality levels and, therefore, do not engage in any avoidance behavior. This is unlikely, as demonstrated via the following correlational relationships between air quality and outdoor time use. Appendix Table A.13 shows the correlations between PM2.5 information and respondents' outdoor time use at baseline and endline. The table shows that a $10\mu\text{g}/\text{m}^3$ increase in the PM2.5 concentration reduces respondents' stated outdoor time use by approximately 3 minutes per day at endline. This negative correlation is absent at baseline, suggesting that respondents adjust their time use on worse air quality days.

Overall, we do not find strong evidence that exposure to an information source alone leads to differential changes in beliefs about the state of air quality or in subsequent avoidance behaviors. Beliefs about air quality levels and the demand for avoidance behavior *per se* may affect the willingness to pay for air quality information. However, our treatments did not have differential effects in these beliefs and are unlikely to have driven differential willingness to pay for air quality information between treatment arms.

6.5 Beliefs about SMS forecast's accuracy

Next, we focus on the recipients' beliefs about the accuracy of information. Note that recipients in the two treatment arms receive identical readings, forecasts, and messages other than the source to which the information is attributed. As such, differential beliefs about the accuracy of the SMS forecasts should be formed through the source to which the information is attributed.

We isolate the respondents' beliefs about the accuracy of the SMS forecasts using outcomes from two incentivized forecast games in the endline survey. We conduct two types of incentivized elicitation regarding air quality forecasts: 1) respondents' belief about the actual air quality level tomorrow and 2) their guess of the SMS forecast. The absolute difference between the two measures captures the respondents' belief about the quality of SMS forecasts conditional on their own beliefs about air quality the next day.

We find that the respondents in the Government arm believe in larger

SMS forecast errors than in the NGO arm. Column 4 on Table 1 shows the difference to be 2.8 points in the concentration measure ($\mu\text{g}/\text{m}^3$). The effect size is 12% of the NGO arm's mean (22.7). The effect is statistically significant at the 5% level and survives adjustments to multiple hypothesis testing.

Those assigned to the Government arm believe, on average, that the SMS forecasts are less accurate than those in the NGO arm despite having statistically indistinguishable willingness to pay for the forecast services. This result suggests that recipients either have a relatively low willingness to pay for forecast precision on the margin or care about other attributes such as punctuality and ease of access. As such, we examine how recipients' beliefs about service quality associated with a particular source are shifted and, more generally, how their preferences between sources shift due to our intervention.

6.6 Preferences for sources

We conduct donation games with financial stakes for our primary measure of respondents' preference between sources. In the baseline and endline surveys, respondents PKR 100 between government and NGO sources, which the survey team donates to respective agencies. The relative allocations, as well as changes to them between baseline and endline, identify respondents' preferences for information sources with real financial stakes.

We find that the respondents shift a larger fraction of their donations to the experimentally assigned sources at endline. Figure 4 shows the distributions for both treatment groups at baseline and endline. The figures show most respondents split the donations 50:50 at baseline, but their preferences diverge significantly by treatment arm at the endline. More than 90 percent of respondents assigned to the Government arm donate more to the government at the endline, as opposed to the NGO one. On the other hand, more than 90 percent of respondents assigned to the NGO arm donate more to the NGO at the endline. The average ratio between the assigned source and the other is approximately 75:25. Column 5 in Table 1 confirms that those assigned to the

Government arm donate PKR 54 more to the government, on average, relative to the respondents in the NGO arm.

Furthermore, we find evidence of a higher willingness to pay for information from the experimentally assigned source when we look within individuals. After the BDM, they are asked hypothetically how much they would be willing to pay if the forecast were to come from the other source (i.e., NGO for those assigned to government, and vice versa). Column 4 in Table 2 shows that the respondents are, on average, willing to pay PKR 16 less for the alternative source than for the experimentally assigned one. Although the hypothetical WTP measure is not based on revealed preference, this finding aligns with other findings from other measures of preferences for sources.

We do not have other revealed-preferences measures that would help us identify exactly what components, besides forecast accuracy, of the information sources' services recipients value. As such, we collect a set of stated-preference measures of the recipients' approval of the sources' reliability (i.e., punctuality of their forecasts), accuracy, and overall service quality. We collect these measures for both Government and NGO service providers for each recipient.

Table 3 shows that recipients assigned to the Government arm have significantly higher approval for it regarding reliability, accuracy, and overall satisfaction than those assigned to the NGO arm. Similarly, Table 4 shows the symmetrical results for those assigned to the NGO arm. We find that the recipients are satisfied with the services they received regardless of treatment arms, and they associate their satisfaction with the source made salient by our intervention. Overall, we find evidence that preferences are relatively malleable in a highly frictional market for information services.

6.7 Heterogeneity by baseline beliefs

6.7.1 Pre-specified dimensions of heterogeneity

So far, we have identified average intend-to-treat effects of information sources on beliefs and preferences. However, one may be concerned that the evolution of beliefs and preferences may depend on consumers’ priors. Average effects may mask significant heterogeneous treatment effects, leading to divergence in beliefs and preferences based on consumers’ baseline characteristics.

To address such concerns, we pre-specify four dimensions to test potential heterogeneity, as described in detail in Appendix Section D.1. They are: a) a measure of relative preferences between the sources using the baseline donation game, b) a relative measure overall approval between sources, c) a relative measure of beliefs about the sources’ accuracy, and d) beliefs about air quality levels. All of these measures are collected at baseline.

We focus the beliefs and preferences between information sources at baseline based on the emerging body of work on media bias, trust for information sources, and polarization. Previous work has shown that agents may place heavier weights on information from a source that aligns with their priors, leading to polarization in preferences and beliefs (e.g., Gentzkow et al. 2023; Chopra et al. 2022).¹³ If, on the other hand, agents do not exhibit belief confirmation or do not hold strong priors about the sources’ quality, they may shift their priors more strongly to information from a source that they are less exposed to at baseline. As such, it is *a priori* unclear how the demand for the sources evolves based on their baseline preferences and beliefs.

We focus on baseline beliefs about air quality to see if the extent of belief-updating depends on the accuracy of baseline beliefs and the recipients’ beliefs about signal quality. Those less well-informed about air quality levels may hold priors with more deviations from the truth. Those individuals may, therefore,

¹³This may be driven by “belief confirmation,” i.e., they prefer sources that distort information toward their prior beliefs (Mullainathan and Shleifer 2005), or driven by uncertainty about accuracy of information sources, inducing an individual to put heavier weights on their preferred source (Gentzkow and Shapiro 2006).

update their beliefs more strongly toward the truth based on the signals they receive and value the SMS forecasts more. Critically, the extent to which such individuals update their beliefs would also depend on their beliefs about the strength of the signal, which they may glean from the information source.

6.7.2 Heterogeneous treatment effects

We estimate heterogeneous treatment effects based on the pre-specified dimensions.¹⁴ The empirical specifications are described in Appendix Section D.2. We do not find strong evidence that the consumers respond differentially based on their prior beliefs about the information sources' service quality. We also find evidence that consumers with higher baseline forecast errors have higher endline forecast errors if they are assigned to the government arm.

First, we do not find strong evidence of heterogeneous treatment effects based on consumers' preferences for sources, except for endline donations to the government. Appendix Tables A.16 to A.18 show the linear heterogeneous treatment effect estimates and their categorical equivalents in Appendix Tables A.20 to A.22. Coefficients on interaction terms from Appendix Tables A.16 to A.17 are not generally statistically significant. One exception is the negative interaction terms for the endline donation outcome (Column 5), which is likely because the outcome measure has a ceiling at PKR 100. In other words, those who report to prefer the government in baseline would donate more to the government and would not be able to increase donations to the government beyond PKR 100. One exception is the marginally significant interaction term in Column 3, Appendix Table A.18, but this result is not corroborated with a categorical specification in Appendix Table A.22.

Second, we find an adverse heterogeneous treatment effect based on the consumers' baseline forecast error on forecast error and respondents' beliefs about the SMS's error Appendix Table A.19 shows the linear estimates, and Appendix Table A.23 the categorical equivalent. For those assigned to the gov-

¹⁴We do not, however, adjust for multiple testing in these secondary hypotheses.

ernment arm relative to the NGO one, having a 100% larger baseline forecast error is associated with having 26% higher endline forecast errors. In other words, those with higher baseline errors update their priors less about air pollution levels than similar individuals if assigned to the government arm v.s. the NGO. Similar causal effects also exist on the respondents' beliefs about the SMS's errors but are less precisely estimated.

Two takeaways emerge as from the pre-specified analysis. First, there are no strong heterogeneous effects based on consumers' priors about the sources on the demand for air quality information. This confirms our results from Section 6.6 that the consumers have relatively weak priors about information sources, and their beliefs are relatively malleable. Second, even when attributions to information sources do not meaningfully affect the demand for the ultimate service (air quality information), consumers with less accurate beliefs about air quality update their beliefs more slowly when they are assigned to the government source, which they believe to have lower quality.

7 Mapping results to a conceptual framework

Our empirical results show that, although the source does not differentially affect the recipients' demand for air quality information, it affects several underlying beliefs and preferences. To put further structure to our findings, we map our empirical results to a conceptual framework that specifies the utility of consuming an air quality forecast service in terms of its attributes. We establish links between pre-specified hypotheses and the conceptual framework. We then identify which attributes of the utility function shift in response to an information source that is exogenously made salient. We present the highlights in this section and provide further details in Appendix Section E.

We specify a consumer's utility function consisting of attributes underpinned by their beliefs. First, consumers hold beliefs about the state variable, i.e., the air quality level. They value the forecast information because it provides additional signals about the state and helps them take better mitigation

measures against air pollution. Second, they also hold beliefs about signal quality (i.e., the accuracy of the SMS forecast) for a given source. Third, they may also hold beliefs and preferences about a source that is not tied to signal quality (e.g., consistent availability) as well as innate preferences for a given source, all of which we bundle into an attribute. Such beliefs and preferences factor into the utility function as attributes. In our experiment, we elicit consumers’ utility gains from a signal—the SMS air quality forecast—whose source we exogenously vary. We denote the utility gained from accessing an air quality forecast from source $s \in \{G, P\}$ of a consumer i assigned to treatment arm $a \in \{G, P\}$ (G(overnment) or P(riate), i.e., NGO) on day t as $u_{i,a,t}^s$. We measure $u_{i,G,t}^G$ and $u_{i,P,t}^P$ through BDM, and $u_{i,G,t}^P$ and $u_{i,P,t}^G$ through hypothetical willingness-to-pay measures.

Based on the utility and other belief measures, our empirical analysis tests whether treatment assignment differentially affects the demand for air quality information and its attributes, i.e., whether being assigned G affects $u_{i,G,t}^G$ and its attributes differently from being assigned to P affecting $u_{i,P,t}^P$. We also analyze whether treated individuals prefer to receive forecasts from the source to which they are experimentally assigned against the alternative, i.e., if $u_{i,G,t}^G > u_{i,G,t}^P$ and $u_{i,P,t}^P > u_{i,P,t}^G$ and changes in beliefs about which attributes explain the difference.

Our empirical results map to the following descriptions of the model’s terms. First, we fail to reject $u_{i,G,t}^G \neq u_{i,P,t}^P$ but find that $u_{i,G,t}^G > u_{i,G,t}^P$ and $u_{i,P,t}^P > u_{i,P,t}^G$, i.e., the willingness to pay is not differentially affected between treatment arms but recipients prefer the source to which they are assigned. Second, we do not find that their air quality forecast error is different between treatment arms, i.e., their beliefs about the state variable are not differentially affected. Third, recipients assigned to the government arm believe that the SMS forecasts have higher errors than those in the NGO arm.

Our empirical results suggest that consumers put a relatively small weight on the signal quality attribute and/or that they value attributes specific to the assigned source other than signal quality. On these potential mechanisms,

we provide correlational evidence in Appendix Section E. First, we show that $u_{i,G,t}^G$ is not correlated with the SMS forecast error, but $u_{i,P,t}^P$ is, suggesting that consumers receiving government forecasts do not value signal quality on the margin. Second, we evaluate what aspects of the service quality, as measured through Likert-scale statements, correlate with $(u_{i,G,t}^G - u_{i,G,t}^P)$ and $(u_{i,P,t}^P - u_{i,P,t}^G)$. We find that it is their approval of the assigned source in terms of reliability. These correlational results suggest that consumers do not value the forecast’s precision at the margin but rather value the reliability of the service that they have experienced.

8 Conclusion

We study how consumers form beliefs and preferences for a public good, an environmental information service, when there is uncertainty about the state of environmental quality as well as the service quality of suppliers. We conduct a randomized control trial in which we randomly attribute air quality forecast services to one of two sources: the government and an NGO. We evaluate if the random attribution leads to a differential demand for air quality forecasts and beliefs about air quality levels and information sources. We also investigate whether respondents hold varying beliefs about the information’s accuracy or exhibit preferences between information sources.

We find that consumers in working-class neighborhoods of Lahore have high demand for air quality information, yet not differentially so by the associated source. Yet, we find that those assigned to the government arm believe the information they receive is less accurate than those in the NGO arm, while they are equally willing to pay for the information service. Consumers shift their preference toward the source to which they are exposed. Our findings provide insights into a market for a public good with competing suppliers and high levels of friction to access them; consumers have limited access to air quality information at baseline yet have a high demand for it as measured through an incentivized elicitation. They prefer a source to which they are

exposed, yet may not value the accuracy of the information at the margin, at least as long as they generally approve of the service quality.

Our results have policy implications for governments, multilateral organizations, and civil society on improving access to environmental information.

First, our findings suggest sizable potential welfare gains from increasing access to environmental information where there are damages from environmental degradation. Residents of a working-class neighborhood of Lahore are willing to pay PKR 119 (\approx USD 0.50 as of August 2022) per month for air pollution forecasts after the conclusion of the free information intervention. This amount roughly equates to 50% of the cost of monthly prepaid mobile and data services. Scaling the service across the city—with close to 14 million residents—will likely lead to a large benefit based only on individual consumers' willingness to pay for information. In contrast, we find in the baseline survey that only around 9 percent of the working-class citizens in Lahore report having accessed air quality readings from a given source, underlying the market failure in the provision of public goods.

Second, our findings show limited sensitivity of the overall willingness to pay to the source and to their believed accuracy, suggesting that the welfare gains would come from access to information regardless of the source. A social planner, therefore, should increase citizens' access to air quality information regardless of the source. Note, however, that such takeaway may be conditional on a certain signal and service quality we maintained as part of the experiment. Potential concerns remain as to whether, outside the experimental setting, there are trade-offs between accuracy, reliability, and operational costs between information sources that have different hardware and other technical capacities. Government agencies tend to have higher quality monitors that can monitor multiple scheduled pollutants and meet technical standards, while crowd-funded NGO need to rely on low-cost monitors of uncertain quality (US EPA 2024). Consumers may be harmed if, for instance, NGO air quality readings are significantly less accurate than the government ones to the point where consumers would care about the difference in accuracy.

We provide two pieces of evidence that, at least in Lahore’s context, such trade-offs are unlikely to be of concern. First, we find that non-government sources perform better across several metrics of accuracy and reliability than government ones. Appendix Tables A.2 to A.4 show that the government (EPD) air quality readings have higher noise than the NGO (PAQI) ones based on low-cost monitors when we define the truth to be the third-party U.S. Consulate measures. Second, when it comes to predictive modeling, forecasts with inputs from government readings perform as well as those using non-government inputs, as shown in Appendix Table A.24 in terms of forecast errors. In other words, there is limited scope for improving air quality forecasts for the city of Lahore by introducing additional devices or sources.

Third, our findings show shifts in consumers’ preferences for the experimentally assigned source against alternatives from a baseline of seemingly weak priors. However, our experimental evidence is unable to speak to the long-term effects on consumers’ preferences for a source. For instance, consumers’ preferences for an information source may become less malleable as they experience services from it. Such patterns could lead to polarized beliefs about the sources’ accuracy and service quality in the long run. Shifting consumers to receive information from one source to another may have a negative welfare impact, as we find that consumers have a lower willingness to pay for information from a less familiar source. Further work is needed to understand consumers’ longer-term belief updating process and preferences.

References

- Acemoglu, D., A. Cheema, A. I. Khwaja, and J. A. Robinson (2020). Trust in state and nonstate actors: Evidence from dispute resolution in pakistan. *Journal of Political Economy* 128(8), 3090–3147.
- Ahmad, H. F., M. Gibson, F. Nadeem, S. Nasim, and A. Rezaee (2022). Forecasts: Consumption, Production, and Behavioral Responses.
- Akerlof, G. A. and R. E. Kranton (2000). Economics and identity. *The quarterly journal of economics* 115(3), 715–753.
- Andrabi, T., J. Das, and A. I. Khwaja (2017). Report cards: The impact of providing school and child test scores on educational markets. *American Economic Review* 107(6), 1535–1563.
- Athey, S. and G. Imbens (2017). The Econometrics of Randomized Experiments. In *Handbook of Economic Field Experiments*, Volume 1, pp. 73–140. Elsevier.
- Barwick, P. J., S. Li, L. Lin, and E. Y. Zou (2024). From fog to smog: The value of pollution information. *American Economic Review* 114(5), 1338–1381.
- Baysan, C. (2022). Persistent Polarizing Effects of Persuasion: Experimental Evidence from Turkey. *American Economic Review* 112(11), 3528–3546.
- Becker, G. M., M. H. Degroot, and J. Marschak (1964). Measuring utility by a single-response sequential method. *Behavioral Science* 9(3), 226–232.
- Belloni, A., V. Chernozhukov, and C. Hansen (2014, April). Inference on Treatment Effects after Selection among High-Dimensional Controls. *The Review of Economic Studies* 81(2), 608–650.
- Bergstrom, T., L. Blume, and H. Varian (1986). On the private provision of public goods. *Journal of Public Economics* 29(1), 25–49.


- Bowen, T. R., D. Dmitriev, and S. Galperti (2023). Learning from shared news: When abundant information leads to belief polarization. *The Quarterly Journal of Economics* 138(2), 955–1000.
- Chopra, F., I. Haaland, and C. Roth (2022). The demand for news: Accuracy concerns versus belief confirmation motives.
- Chowdhury, A., T. Garg, M. Jagnani, and M. Mattsson (2024). Adoption, use, and effects of air purifiers in households. *Available at SSRN 4665018*.
- Coury, M., L. Falconer, and A. La Nauze (2024). Wildfire smoke and private provision of public air-quality monitoring.
- Das, J., A. Holla, A. Mohpal, and K. Muralidharan (2016, December). Quality and Accountability in Health Care Delivery: Audit-Study Evidence from Primary Care in India. *American Economic Review* 106(12), 3765–3799.
- Davies, C. (2023, 4 May). Polarised politics are tearing Pakistan apart. *BBC*.
- Dhinakar Bala, A., M. Mattsson, and S. Vyas (2024). Why do governments make promises they cannot keep? evidence from india’s national clean air programme. Technical report, SSRN.
- Freeman, R., W. Liang, R. Song, and C. Timmins (2019, March). Willingness to pay for clean air in China. *Journal of Environmental Economics and Management* 94, 188–216.
- Gentzkow, M. and J. Shapiro (2006). Media Bias and Reputation. *Journal of Political Economy* 114(2), 280–316. Publisher: The University of Chicago Press.
- Gentzkow, M., M. B. Wong, and A. T. Zhang (2023). Ideological Bias and Trust in Information Sources.
- Ghanem, D. and J. Zhang (2014, September). ‘Effortless Perfection:’ Do Chinese cities manipulate air pollution data? *Journal of Environmental Economics and Management* 68(2), 203–225.

- Greenstone, M., T. Ganguly, C. Hasenkopf, N. Sharma, and G. Hrishikesh (2024). Annual update. *Air Quality Life Index*.
- Greenstone, M., K. Lee, and H. Sahai (2021). Indoor air quality, information, and socioeconomic status: Evidence from delhi. In *AEA Papers and Proceedings*, Volume 111, pp. 420–424. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- GSMA (2024). The Mobile Gender Gap Report.
- Hirshleifer, S., M. Naseem, A. A. Raza, and A. Rezaee (2023). The spread of (mis) information: A social media experiment in pakistan.
- Hsiang, S., P. Oliva, and R. Walker (2019). The distribution of environmental damages. *Review of Environmental Economics and Policy*.
- Ito, K. and S. Zhang (2020, May). Willingness to Pay for Clean Air: Evidence from Air Purifier Markets in China. *Journal of Political Economy* 128(5), 1627–1672. Publisher: The University of Chicago Press.
- Jha, A. and A. La Nauze (2022, November). US Embassy air-quality tweets led to global health benefits. *Proceedings of the National Academy of Sciences* 119(44), e2201092119. Publisher: Proceedings of the National Academy of Sciences.
- Khan, A., S. Nasim, M. Shaukat, and A. Stegmann (2021). Building trust in the state with information: Evidence from urban punjab. *Journal of Public Economics* 202, 104494.
- Lan, R., S. D. Eastham, T. Liu, L. K. Norford, and S. R. Barrett (2022). Air quality impacts of crop residue burning in india and mitigation alternatives. *Nature communications* 13(1), 6537.
- Mu, Y., E. Rubin, and E. Y. Zou (2024). What’s missing in environmental self-monitoring: Evidence from strategic shutdowns of pollution monitors. *Review of Economics and Statistics*, 1–45.

- Mullainathan, S. and A. Shleifer (2005, September). The Market for News. *American Economic Review* 95(4), 1031–1053.
- Muralidharan, K. and V. Sundararaman (2015, August). The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India. *The Quarterly Journal of Economics* 130(3), 1011–1066.
- Raza, A. (2021, 17 Nov). Punjab minister wants FIA action against fake air quality data. *Financial Times*.
- Samuelson, P. A. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics* 36(4), 387–389.
- Shapiro, C. and H. Varian (1999). *Information rules: A strategic guide to the network economy*. Harvard Business School Press.
- The Parliament of Pakistan (1997). Pakistan Environmental Protection Act.
- US EPA (2024). Low-Cost Air Pollution Monitors and Indoor Air Quality. Technical report.
- White, J. E. (2018). AirNow Global – Platform for International Collaboration.
- WHO (2021). *WHO global air quality guidelines: particulate matter (PM_{2.5} and PM₁₀), ozone, nitrogen dioxide, sulfur dioxide and carbon monoxide*. World Health Organization.
- Zivin, J. S. G., B. Krebs, and M. J. Neidell (2024). Private adoption of public good technologies: The case of purpleair. Technical report, National Bureau of Economic Research.

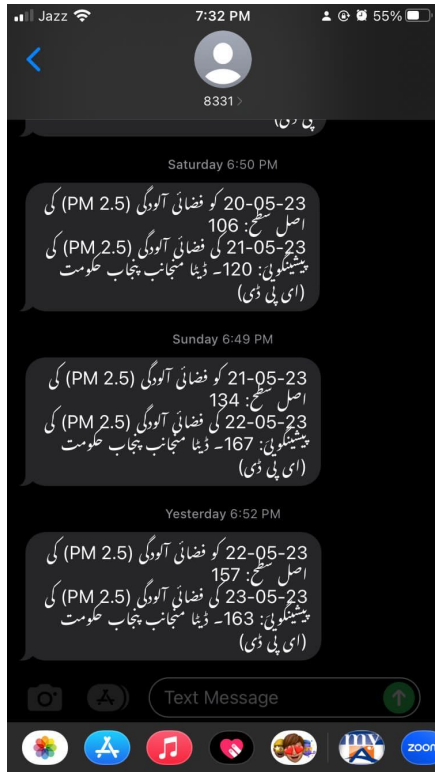
9 Figures

Figure 1: Air quality readings from three main sources



This figure shows the daily average PM2.5 concentration (in $\mu\text{g}/\text{m}^3$) levels by sources. “American consulate” refers to readings from the air quality monitor at the American consulate in Lahore. We treat this reading as the ground truth. “PAQI” refers to readings from the average of lower-cost air quality monitors managed by Pakistan Air Quality Initiative (PAQI), an NGO based in Lahore. “EPD” refers to readings from air quality monitors managed by the Environmental Protection Department (EPD) of the Government of Punjab Province.

Figure 2: Sample messages to respondents



(a) T1: Government (EPD)



(b) T2: NGO (PAQI)

The figures above are screenshots from research managers' cellphones showing daily messages from T1 (EPD) and T2 (PAQI). Daily messages are delivered from the same short code (8331) so that recipients can compare daily readings and forecasts. For 20 May, 2023, the daily messages read as follows:


T1: Actual Air Quality (PM 2.5) on 20-05-23: 106

Air Quality Forecast (PM 2.5) for 21-05-23 using data From Punjab Government (EPD): 120

T2: Actual Air Quality (PM 2.5) on 20-05-23: 106


Air Quality Forecast (PM 2.5) for 21-05-23 using data From NGO (PAQI): 120

Figure 3: Demand curves for air pollution forecast by treatment




This figure shows the distributions of respondents' bids for two months of air pollution forecast service from the endline survey. "Government" corresponds to T1, the arm in which the EPD source is made salient. "NGO" corresponds to T2, the arm in which the PAQI source is made salient.

Figure 4: Baseline and endline donation to government sources vs NGO



(a) Baseline



(b) Endline

This figure shows the distributions of respondents' donations to a government agency vs. a non-government entity for environmental protection, measured at the baseline and endline surveys. The measure is defined as the amount out of PKR 100 donated to the government source. "Government" corresponds to T1, the arm in which the EPD source is made salient. "NGO" corresponds to T2, the arm in which the PAQI source is made salient.

10 Tables

Table 1: Pre-specified hypotheses: ITT

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
Constant	237.5*** (2.19)				
Gov't arm		0.33 (3.68)	0.051 (0.040)	2.82** (1.29)	53.8*** (1.04)
P value	0	.927	.208	.029	0
Q value	.001	.351	.116	.03	.001
Observations	993	993	993	991	989
Endline mean of NGO		237.2	0.73	22.7	22.9

Notes: Model: PDSLASSO. “WTP”: Willingness to pay for two months of SMS air quality forecasts. “Forecast error”: the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. “SMS error”: the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. “Donation gov’t”: amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$. We also show the critical values for the “Constant” and “Gov’t arm” coefficients. “P value:” Unadjusted p-values. “Q value”: Benjamini Krieger Yekutieli (2006) sharpened q-values.

Table 2: ITT: Alternative definitions of the WTP outcome

	(1)	(2)	(3)	(4)
	WTP	WTP (other)	diff(WTP)	diff(WTP)
Gov't arm	0.33 (3.68)	0.55 (3.66)	-0.21 (0.54)	
Constant				15.9*** (0.60)
Observations	993	993	993	993
Endline mean of NGO	237.2	221.2	16.0	

Notes: “WTP”: The pre-specified outcome measuring the willingness to pay for two months of SMS air quality forecasts, where the assigned source is made salient. “WTP if other source”: hypothetical WTP if the forecast were to come from the other source not assigned to them. “diff(WTP sources)”: the difference between the willingness to pay for the assigned vs. the other sources. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table 3: ITT: Stated preference measure on satisfaction with the Government service

	(1)	(2)	(3)	(4)
	Gov't: Approval index	Gov't: Reliable	Gov't: Accurate	Gov't: Approve
Gov't arm	2.25*** (0.054)	1.41*** (0.037)	1.34*** (0.039)	1.33*** (0.037)
Observations	989	980	950	989
Endline mean of NGO	-2.27	3.07	3.05	3.10

Notes: We present estimates of effects on the stated-preference measures on the respondents' satisfaction with the Government's service. We ask if they are overall satisfied with the service (Column 3), if they think the service is reliable and on time (Column 4), and if they believe the forecasts are accurate (Column 2), in the Likert scale where positive values indicate approval. The measure for Column 1 is a standardized sum of measures in Columns 2 and 4. The measures for Columns 2 through 4 are in the 5-point Likert scale. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table 4: ITT: Stated preference measure on satisfaction with the NGO service

	(1)	(2)	(3)	(4)
	NGO: Approval index	NGO: Reliable	NGO: Accurate	NGO: Approve
Gov't arm	-1.69*** (0.039)	-1.34*** (0.035)	-1.49*** (0.037)	-1.37*** (0.033)
Observations	982	972	958	986
Endline mean of NGO	1.70	4.46	4.47	4.42

Notes: We present estimates of effects on the stated-preference measures on the respondents' satisfaction with the NGO source's service. We ask if they are overall satisfied with the service (Column 3), if they think the service is reliable and on time (Column 4), and if they believe the forecasts are accurate (Column 2), in the Likert scale where positive values indicate approval. The measure for Column 1 is a standardized sum of measures in Columns 2 and 4. The measures for Columns 2 through 4 are in the 5-point Likert scale. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

A Appendix tables and figures

Table A.1: Balance table of outcomes and key demographic variables at baseline

Variable	(1) NGO Mean/(SE)	(2) Government Mean/(SE)	(1)-(2) Pairwise t-test Mean difference
Baseline forecast error	0.725 (0.019)	0.714 (0.019)	0.011
Baseline donation gov't	50.139 (0.682)	50.119 (0.654)	0.020
Baseline: hours spent outside	7.403 (0.204)	7.440 (0.198)	-0.037
Stated preference for citizens group	0.013 (0.042)	-0.011 (0.043)	0.024
Stated preference for government	-0.009 (0.043)	-0.010 (0.043)	0.001
Comprehended the text message without explanation	0.768 (0.019)	0.766 (0.019)	0.002
Received air pollution info from: EPD	0.087 (0.013)	0.083 (0.012)	0.004
Received air pollution info from: AirVisual App	0.097 (0.013)	0.089 (0.013)	0.008
Index: Sentiment on air quality	-0.019 (0.032)	0.010 (0.032)	-0.029
Asset index	0.020 (0.046)	-0.026 (0.043)	0.046
F-test of joint significance (F-stat)			0.210
Number of observations	504	504	1008

Notes: This table presents sample means and standard deviations by treatment arms, mean differences and their t-tests, and the two-tailed significance. All measures come from the baseline survey. "Baseline forecast error": baseline measure of the pre-specified forecast-error outcome. "Baseline donation gov't": baseline measure of the preference for the government source vs the NGO. "Baseline: hours spent outside": time spent outdoors, as calculated from a time-use log. "Stated preference for citizens group": indexed measure of respondents' stated beliefs that a) air quality readings from the NGO are accurate, and that b) they approve of the job that the NGO is doing to address air quality. "Stated preference for government": indexed measure of respondents' stated beliefs that a) air quality readings from the government are accurate, and that b) they approve of the job that the government is doing to address air quality. "Comprehended the text message without explanation": When the respondent was shown a mock-up of a text message they will receive, they understood it without further explanation. "Received air pollution info from: EPD": self-reported to have accessed air quality readings from EPD. "Received air pollution info from: Air Visual App": self-reported to have accessed air quality readings from the AirVisual App, on which PAQI disseminates air quality information. "Index: Sentiment on air quality": indexed measure that a) respondents care about air quality in places they live, b) they have been concerned about air quality in general in the last week, c) their quality of life is significantly affected at home, their performance at work or school is significantly affected, d) their sleep is affected, they reduced the number of hours worked, and e) the number of days in the last week with unsatisfactory air quality. "Asset index": indexed measure of the household's ownership of electronic appliances. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table A.2: Summary statistics: PM2.5 readings by source

	(1)		
	mean	sd	count
Pre-intervention			
AirNow	209.0	114.9	107
EPD	171.4	86.7	80
PAQI	161.7	64.3	116
Urban	210.5	108.3	107
Sprintars	69.2	19.5	116
During/post-intervention			
AirNow	63.5	39.4	175
EPD	55.7	36.0	173
PAQI	58.7	31.0	190
Urban	92.4	70.9	134
Sprintars	59.1	14.5	186
Total			
AirNow	118.7	104.6	282
EPD	92.3	78.4	253
PAQI	97.7	68.3	306
Urban	144.9	106.9	241
Sprintars	63.0	17.3	302

Notes: “Pre-intervention”: time period prior to our intervention (Feb 18), i.e., the period with high levels of PM2.5 concentrations. “During/post-intervention”: Period since February 18, when there are relatively low PM2.5 concentrations. “Total”: readings from November 1, 2022 to August 26, 2023. “AirNow”: U.S. Consulate readings. “Urban”: The Urban Unit (Provincial Government of Punjab). “PAQI”: Pakistan Air Quality Initiative. “EPD”: Environment Protection Department (Provincial Government of Punjab). “Sprintars”: Satellite-based measure.

Table A.3: Correlations between readings

	AirNow	EPD	PAQI	Urban	Sprintars
AirNow	1				
EPD	0.61***	1			
PAQI	0.82***	0.70***	1		
Urban	0.76***	0.58***	0.73***	1	
Sprintars	0.20***	0.28***	0.27***	0.23***	1

Notes: Pairwise correlation measures of air quality readings by source. “AirNow”: U.S. Consulate readings. “Urban”: The Urban Unit (Provincial Government of Punjab). “PAQI”: Pakistan Air Quality Initiative. “EPD”: Environment Protection Department (Provincial Government of Punjab). “Sprintars”: Satellite-based measure. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.4: Deviations of monitor readings from the American Consulate readings

	RMSE:All	RMSE:pre	RMSE:during/post	MAD:All	MAD:pre	MAD:during/post
EPD	75.6	121.4	41.2	46.4	87.3	28.1
PAQI	63.7	100.8	23.8	34.6	69.5	14.6
Urban	67.6	62.1	71.5	38.7	41.6	36.5
Sprintars	113.4	177.4	44.2	73.2	143.6	32.1

Notes: Deviation from the American Consulate readings by source. RMSE: Root mean squared error. MAD: mean absolute difference. “All”: readings from November 1, 2022 to August 26, 2023. “pre”: time period prior to our intervention (Feb 18), i.e., the period with high levels of PM2.5 concentrations. “during/post”: Period since February 18, when there are relatively low PM2.5 concentrations. “AirNow”: U.S. Consulate readings. “Urban”: The Urban Unit (Provincial Government of Punjab). “PAQI”: Pakistan Air Quality Initiative. “EPD”: Environment Protection Department (Provincial Government of Punjab). “Sprintars”: Satellite-based measure.

Table A.5: Treatment effects on the distribution of willingness-to-pay in 50-PKR bins

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	1-50	51-100	101-150	151-200	201-250	251-300	301-350	351-400
Gov’t arm	0.0083 (0.0068)	-0.0082 (0.011)	-0.026** (0.013)	0.019 (0.020)	-0.00021 (0.024)	0.0053 (0.024)	0.0072 (0.015)	-0.0060 (0.0090)
Observations	1008	1008	1008	1008	1008	1008	1008	1008
Endline mean of NGO	0.012	0.052	0.079	0.16	0.30	0.26	0.087	0.032

Notes: Model: PDSLASSO. The outcomes are dummy variables that equals 1 if the endline willingness-to-pay (WTP) falls in the bin, and 0 otherwise. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table A.6: Tests of distribution of the willingness-to-pay—Kolmogorov-Smirnov and Wilcoxon rank-sum tests

	(1)	(2)
	Kolmogorov–Smirnov statistic	p-value
K-S NGO is smaller	.041	.44
K-S Government is smaller	-.021	.81
K-S Combined	.041	.807
Wilcoxon rank-sum test	.	.919

Notes: The table reports the test statistics and asymptotic p-values from the two-sided Kolmogorov–Smirnov and Wilcoxon rank-sum tests. The two groups are the Government and NGO arms. The Kolmogorov–Smirnov statistic is the supremum of the differences between the two groups. In a row labeled “K-S NGO is smaller,” we test the hypothesis that the distribution of the willingness to pay is lower for the NGO arm. In a row labeled “K-S Government is smaller,” we test the hypothesis that the distribution of the willingness to pay is lower for the Government arm. “K-S Combined” is the combined test statistic. Wilcoxon rank-sum test (also known as the Mann-Whitney two-sample statistic) is a test where the null hypothesis is that the willingness-to-pay measure for Government and NGO arms is drawn from the same distribution.

Table A.7: Pre-specified hypotheses: ITT (winsorized)

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
Constant	237.4*** (2.18)				
Gov't arm		0.31 (3.67)	0.051 (0.037)	1.89* (1.01)	53.8*** (1.04)
Observations	993	993	993	991	989
Endline mean of NGO		237.2	0.73	22.7	22.9

Notes: We winsorize the outcome variables at the 1st and 99th percentiles. “WTP”: Willingness to pay for two months of SMS air quality forecasts. “Forecast error”: the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. “SMS error”: the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. “Donation gov’t”: amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table A.8: Adjustments for multiple hypothesis testing on pre-specified hypotheses (winsorized)

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
P value	0	.932	.177	.061	0
Q value	.001	.284	.113	.065	.001

Notes: We show the critical values for the “Constant” and “Gov’t arm” coefficients in the corresponding columns of Table A.7. “P value:” Unadjusted p-values. “Q value”: Benjamini Krieger Yekutieli (2006) sharpened q-values.

Table A.9: ITT: Alternative definitions of the forecast outcome

	(1)	(2)	(3)	(4)
	abs(own - truth)/truth	(own - truth)/truth	abs(own - truth)	(own - truth)
Gov't arm	0.051 (0.040)	0.060 (0.049)	-0.48 (2.17)	3.39 (2.85)
Observations	993	993	993	993
Endline mean of NGO	0.73	0.47	40.0	91.9

Notes: We present estimates of effects on forecast outcomes with different definitions, where “own” stands for the respondent’s own forecast of the air quality level the next day, and “truth” the actual readings on the corresponding day. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table A.10: ITT: Concerns about air quality

	(1)	(2)
	Care about AQ	N. days good air
Gov't arm	0.0088 (0.046)	0.025 (0.054)
Observations	992	961
Endline mean of NGO	2.59	3.16

Notes: We present estimates of effects on measures of concern about air quality. “Care about AQ”: a Likert-scale measure of how much the respondent cares about air quality in the places they live and work. “N. Days good air”: Number of days in the last week with acceptable air quality. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.11: Secondary outcomes: Time use

	(1)	(2)	(3)
	Endline: hours spent outside	Endline: Hrs (stated)	Endline: Hrs (if bad day)
Gov't arm	-0.036 (0.15)	0.010 (0.12)	0.0063 (0.11)
Observations	993	993	993
Endline mean of NGO	5.14	3.89	3.65

Notes: Standard errors are reported in parentheses. For the “[endline]: hours spent outside” variable, we ask respondents the type of activity (sleep, paid work, homemaking, leisure, travel, and other) they conducted for each hour of the previous day and whether it was indoors or outdoors. We aggregate the number of hours the respondent engaged in any outdoor activity. For the “[endline]: Hrs (stated)” variable, we ask respondents to state how many hours they spent outside the previous day, as opposed to aggregating the hours using the time-use module. For the “[endline]: Hrs (if bad day)” variable, we ask respondents how many hours they would spend on a bad air quality day. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.12: Secondary outcomes: Mask use

	(1)	(2)	(3)
	Has mask	Shows mask	Uses mask
Gov't arm	-0.035* (0.021)	-0.011 (0.015)	-0.043** (0.021)
Observations	993	993	993
Endline mean of NGO	0.20	0.099	0.19

Notes: Column 1 “Has mask” refers to a binary outcome in which the surveyed individual responded to have purchased or been given a mask. Column 2 indicates whether the individual showed a mask to the enumerator. Column 3 indicates whether the individual reported to have used a mask in the last week. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.13: Correlations: Readings, forecasts, and outdoor time use

	(1)	(2)	(3)	(4)
	Baseline: hours spent outside	Endline: hours spent outside	Endline: hours spent outside	Endline: hours spent outside
PM2.5 reading	-0.00060 (0.0010)			
PM2.5 reading		-0.0054** (0.0023)		-0.0049** (0.0023)
PM2.5 forecast (SMS)			-0.0055 (0.0037)	-0.0045 (0.0037)
Observations	1007	992	992	992

Notes: Standard errors are reported in parentheses. We regress baseline and endline outdoor time use measures on PM2.5 readings and forecasts on the relevant days. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.14: Secondary outcomes: Preference for air quality policies over other domains

	(1)	(2)	(3)
	AQ over Educ	AQ over health	AQ over waste
Gov't arm	0.0038 (0.012)	0.0046 (0.014)	0.021 (0.020)
Observations	992	992	993
Endline mean of NGO	0.050	0.077	0.17

Notes: Standard errors are reported in parentheses. The outcome is defined as 1 if they prefer the government invest in air quality v.s. other policy goals. We ask a hypothetical scenario in which the local government has PKR 100 million to allocate either towards improving air quality or towards investing in one of three other goals (education, health, and waste management for Columns 1, 2, and 3, respectively). All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.15: Secondary outcomes: Demand for filing complaints about air quality

	(1)	(2)
	Takes info	Plans to complain
Gov't arm	-0.016 (0.018)	-0.0024 (0.017)
Observations	993	993
Endline mean of NGO	0.85	0.12

Notes: Standard errors are reported in parentheses. "Takes info:" At the end of the endline survey, we prompt the respondent that EPD is a government agency responsible for addressing air quality issues in Lahore. We tell the respondents that we have a document that shows them how to file a complaint to the EPD and ask if they would like a copy. The outcome is defined as 1 if the respondent takes a pamphlet. "Plans to complain:" The outcome is defined as 1 if a respondent intends to file a complaint to the EPD about air quality. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.16: Heterogeneous effects: Baseline donation to government

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
Baseline donation gov't	0.43** (0.17)	0.83 (0.66)	-0.011** (0.0053)	-0.048 (0.11)	0.50*** (0.13)
Gov't arm		4.39 (16.6)	-0.052 (0.22)	-0.25 (2.81)	95.2*** (3.02)
Gov't arm × Baseline donation gov't		-0.074 (0.30)	0.0020 (0.0040)	0.052 (0.056)	-0.83*** (0.056)
Observations	993	993	993	991	989
Endline mean of NGO		237.2	0.73	22.7	22.9

Notes: Heterogeneous treatment effects by donation to government at baseline. “WTP”: Willingness to pay for two months of SMS air quality forecasts. “Forecast error”: the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. “SMS error”: the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. “Donation gov't”: amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.17: Heterogeneous effects: Baseline overall approval of government v.s. NGO source

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
Relative stated prf for govt: Approval	8.50*** (2.20)	0 (.)	0 (.)	0 (.)	0 (.)
Gov't arm		0.32 (3.68)	0.050 (0.040)	2.73** (1.29)	53.8*** (0.94)
Gov't arm × Relative stated prf for govt: Approval		1.89 (4.17)	-0.020 (0.062)	0.35 (1.05)	-14.5*** (0.79)
Observations	990	990	990	988	986
Endline mean of NGO		237.0	0.73	22.7	23.0

Notes: Heterogeneous treatment effects by a relative measure of overall approval for the government source to the NGO's. The measure “Relative stated prf for govt: Approval” is a standardized difference of Likert-scale questions in which the respondents evaluated their overall approval of the government's and NGO's job in addressing air quality in Lahore. “WTP”: Willingness to pay for two months of SMS air quality forecasts. “Forecast error”: the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. “SMS error”: the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. “Donation gov't”: amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.18: Heterogeneous effects: Baseline belief on information accuracy

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
Relative stated prf for govt: Accuracy	8.03*** (2.17)	0 (.)	0 (.)	0 (.)	0 (.)
Gov't arm		-0.40 (3.60)	0.069* (0.039)	3.14** (1.51)	51.3*** (0.96)
Gov't arm × Relative stated prf for govt: Accuracy		-1.04 (3.76)	-0.070* (0.039)	-0.55 (1.29)	-13.7*** (0.95)
Observations	948	948	948	947	945
Endline mean of NGO		236.4	0.71	23.4	23.8

Notes: Heterogeneous treatment effects by a relative measure of beliefs about the accuracy of the government source's and the NGO's air quality readings. The measure “Relative stated prf for govt: Accuracy” is a standardized difference of Likert-scale questions in which the respondents evaluated their how accurate the government's and NGO's air quality readings are. “WTP”: Willingness to pay for two months of SMS air quality forecasts. “Forecast error”: the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. “SMS error”: the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. “Donation gov't”: amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.19: Heterogeneous effects: Baseline forecast error

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
Baseline forecast error	-1.90 (5.15)	32.0 (26.7)	2.14*** (0.30)	7.49 (9.11)	31.4*** (7.31)
Gov't arm		-6.57 (8.12)	-0.14* (0.084)	-1.69 (2.38)	63.8*** (2.10)
Gov't arm \times Baseline forecast error		10.6 (8.94)	0.26** (0.11)	6.66* (3.77)	-13.8*** (2.39)
Observations	993	993	993	991	989
Endline mean of NGO		237.2	0.73	22.7	22.9

Notes: Heterogeneous treatment effects by the baseline forecast error. "Baseline forecast error" the baseline outcome measure of respondents' forecast error. "WTP": Willingness to pay for two months of SMS air quality forecasts. "Forecast error": the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. "SMS error": the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. "Donation gov't": amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.20: Heterogeneous effects: Baseline donation to government (as categories)

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
More to NGO	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
50-50	3.23 (7.57)	-1.85 (13.3)	-0.0078 (0.11)	0.71 (2.89)	10.5*** (2.94)
More to Govt	13.8* (7.54)	-2.18 (16.5)	-0.13 (0.14)	-0.42 (3.88)	25.1*** (3.52)
Gov't arm		-10.2 (12.6)	0.015 (0.13)	1.13 (1.95)	72.4*** (2.52)
More to NGO \times Gov't arm		0 (.)	0 (.)	0 (.)	0 (.)
50-50 \times Gov't arm		12.9 (13.6)	0.042 (0.15)	1.59 (2.71)	-15.3*** (3.03)
More to Govt \times Gov't arm		12.1 (14.0)	0.0027 (0.15)	1.39 (3.01)	-37.6*** (2.87)
Observations	993	993	993	991	989
Endline mean of NGO		237.2	0.73	22.7	22.9

Notes: Heterogeneous treatment effects by donation to government at baseline. "WTP": Willingness to pay for two months of SMS air quality forecasts. "Forecast error": the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. "SMS error": the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. "Donation gov't": amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.21: Heterogeneous effects: Baseline overall approval of information sources (as categories)

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
Approval: Neutral	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
Approval: Govt	23.2*** (4.60)	2.72 (10.4)	0.056 (0.100)	6.04* (3.12)	21.7*** (2.56)
Approval: Citizen	11.8 (7.68)	22.7* (13.0)	0.049 (0.13)	-4.09 (3.32)	5.78* (3.33)
Gov't arm		1.76 (7.17)	0.16** (0.070)	3.91* (2.36)	73.1*** (1.63)
Approval: Neutral \times Gov't arm		0 (.)	0 (.)	0 (.)	0 (.)
Approval: Govt \times Gov't arm		1.02 (8.54)	-0.16* (0.091)	-1.94 (3.08)	-39.6*** (2.00)
Approval: Citizen \times Gov't arm		-13.0 (14.0)	-0.25 (0.16)	-3.67 (3.44)	-13.7*** (3.28)
Observations	993	993	993	991	989
Endline mean of NGO		237.2	0.73	22.7	22.9

Notes: Heterogeneous treatment effects by a relative measure of overall approval for the government source to the NGO's. The measure is based on Likert-scale questions in which the respondents evaluated their overall approval of the government's and NGO's job in addressing air quality in Lahore. "Approval: Neutral": approves of government as much as the NGO. "Approval: Govt": approves of the government more than the NGO. "Approval: Citizen": approves of the NGO more than the government. "WTP": Willingness to pay for two months of SMS air quality forecasts. "Forecast error": the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. "SMS error": the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. "Donation gov't": amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.22: Heterogeneous effects: Baseline belief on information accuracy (as categories)

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
Accuracy: Neutral	0 (.)	0 (.)	0 (.)	0 (.)	0 (.)
Accuracy: Govt	24.0*** (5.08)	8.88 (10.8)	-0.29** (0.13)	2.14 (3.04)	23.1*** (2.82)
Accuracy: Citizen	17.5** (7.78)	30.2** (12.5)	0.059 (0.15)	6.26** (3.16)	11.9*** (3.41)
Gov't arm		12.0 (7.76)	0.10 (0.078)	5.98** (2.77)	77.4*** (1.70)
Accuracy: Neutral × Gov't arm		0 (.)	0 (.)	0 (.)	0 (.)
Accuracy: Govt × Gov't arm		-17.1* (9.09)	-0.091 (0.098)	-4.95 (3.39)	-40.7*** (2.06)
Accuracy: Citizen × Gov't arm		-13.5 (14.0)	-0.049 (0.14)	-7.50* (3.86)	-21.0*** (3.19)
Observations	993	993	993	991	989
Endline mean of NGO		237.2	0.73	22.7	22.9

Notes: Heterogeneous treatment effects by a relative measure of beliefs about the accuracy of the government source's and the NGO's air quality readings. The measure is based on Likert-scale questions in which the respondents evaluated their how accurate the government's and NGO's air quality readings are. "Accuracy: Neutral": believes government is as accurate as the NGO. "Accuracy: Govt": believes that the government is more accurate than the NGO. "Accuracy: Citizen": believes that the NGO is more accurate than the government. "WTP": Willingness to pay for two months of SMS air quality forecasts. "Forecast error": the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. "SMS error": the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. "Donation gov't": amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: p<0.1*; p<0.05**; p<0.01***.

Table A.23: Heterogeneous effects: Baseline forecast error (as categories)

	(1)	(2)	(3)	(4)	(5)
	WTP	WTP	Forecast error	SMS error	Donation gov't
Baseline error below median	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
Baseline error at or above median	-1.12	11.4	-0.18*	-9.29***	7.21***
	(4.39)	(9.36)	(0.096)	(2.41)	(2.51)
Gov't arm		-1.29	-0.062	0.11	60.0***
		(6.09)	(0.061)	(1.44)	(1.56)
Baseline error below median × Gov't arm		0	0	0	0
		(.)	(.)	(.)	(.)
Baseline error at or above median × Gov't arm		2.20	0.22**	5.58*	-12.2***
		(7.93)	(0.088)	(2.88)	(2.19)
Observations	993	993	993	991	989
Endline mean of NGO		237.2	0.73	22.7	22.9

Notes: Heterogeneous treatment effects by baseline forecast error. “Baseline error below median”: their baseline error is lower than the median. “Baseline error at or above median”: their baseline error is at or higher than the median. “WTP”: Willingness to pay for two months of SMS air quality forecasts. “Forecast error”: the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. “SMS error”: the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. “Donation gov't”: amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table A.24: Forecast errors of the individual and ensemble models

	RMSE:All	RMSE:pre	RMSE:during/post	MAD:All	MAD:pre	MAD:during/post
EPD	68.2	103	39.4	47.3	78.6	31
PAQI	67.1	101.7	39.1	46	76.9	30.3
AirNow	66.3	99.4	39.9	46	74.3	31.6
Urban	67.5	101.7	40.1	46.3	76.7	30.8
Sprintars	115	178.6	48.5	75.8	145.7	35
Ensemble (model in SMS)	77.8	111.2	52.7	55.1	83.1	40.5
Ensemble (only using EPD & PAQI)	67.6	102.2	38.9	46.5	77.7	30.3

Notes: Forecast errors against the American Consulate readings by source. RMSE: Root mean squared error. MAD: mean absolute difference. “All”: readings from November 1, 2022 to August 26, 2023. “pre”: time period prior to our intervention (Feb 18), i.e., the period with high levels of PM2.5 concentrations. “during/post”: Period since February 18, when there are relatively low PM2.5 concentrations. “AirNow”: U.S. Consulate readings. “Urban”: The Urban Unit (Provincial Government of Punjab). “PAQI”: Pakistan Air Quality Initiative. “EPD”: Environment Protection Department (Provincial Government of Punjab). “Sprintars”: Satellite-based measure. “Ensemble (model in SMS)” is constructed based on a weighted average of all sources listed. “Ensemble (only using EPD & PAQI)” is constructed based on a weighted average of EPD and PAQI forecasts.

Figure A.1: Example of the daily EPD report (December 17, 2022)


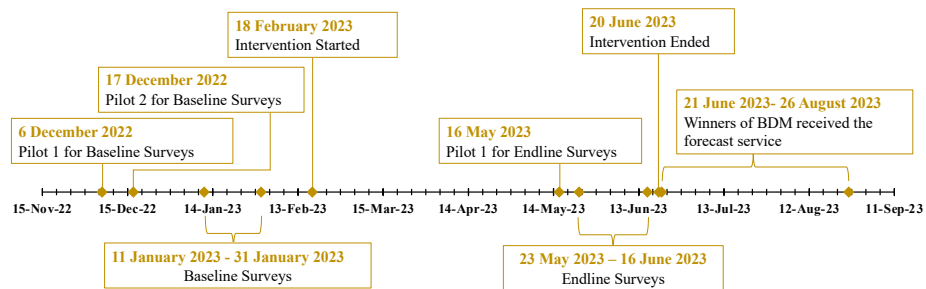


Figure A.2: Timeline of intervention and surveys



B Additional implementation details

B.1 Additional SMS texts

B.1.1 Introductory message

The following messages were sent to the subjects shortly before the beginning of the SMS intervention, which started on 18 February, 2023.

- T1: “Assalam u alaikum! We visited your residence last month and did a survey on Air Pollution in Lahore where you agreed to receive air quality forecast information messages. You will be receiving these messages every day for the next 2 months.

These messages are based on PM 2.5 data which is measured in micrograms per meter cube. The data is collected from the Punjab government’s Environmental Protection Department (EPD) which is tasked with collecting information on Air Pollution. If you have any queries or questions about these messages, please contact the following number [telephone number].”

- T2: “Assalam u alaikum! We visited your residence last month and did a survey on Air Pollution in Lahore where you agreed to receive air quality forecast information messages. You will be receiving these messages every day for the next 2 months.

These messages are based on PM 2.5 data which is measured in micrograms per meter cube. The data is collected from a non-governmental organization (NGO) called Pakistan Air Quality Initiative (PAQI) which collects data on air pollution. If you have any queries or questions about these messages, please contact the following number [telephone number].”¹⁵

¹⁵We use the shorthand “NGO” to refer to organizations of a type, such as PAQI, for the purpose of familiarity with our subjects.

B.1.2 Fortnightly reminder messages

Starting on Saturday, 4 March 2023, reminder messages are sent every two weeks on Saturday about the source and the unit of measurement. The messages by the treatment groups are as follows:

- T1: “The following messages on air pollution (PM 2.5) are based on data from the Punjab Governments Environment Protection Department (EPD). The data is measured in micrograms per meter cube.”
- T2: “The following messages on air pollution (PM 2.5) are based on data from a non-government organization (NGO) named Pakistan Air Quality Initiative (PAQI). The data is measured in micrograms per meter cube.”

B.2 Survey data

B.2.1 Survey frequency

We conduct the following surveys:

- Baseline survey (11th to 31st January 2023)
- Endline survey (29th May to mid/late June 2023)

B.2.2 Survey modules

In the baseline survey, we ask for demographics, some of the outcome measures (i.e., outcomes that are not contingent on the subjects’ having experienced the forecast service), and dimensions of heterogeneity. Detailed survey instruments are included in the appendix. We provide detailed descriptions of outcomes and other variable definitions in Section 4.

The baseline survey modules are as follows:

- Identification of a decision maker in the household as the respondent and consent

- Household roster and their demographics
- Awareness about air pollution in Lahore and access to information
- Donation game between EPD and PAQI, and stated preferences for the sources
- Stated beliefs in their trust in government services
- Incentivized forecast of air pollution (PM 2.5) concentration tomorrow
- Attitudes and behaviors regarding air pollution
- Time use survey and outdoor activities
- Participation in the local community and civil society
- Access to news sources and preferred channels
- Household assets

The endline survey modules are as follows:

- Identification of the same respondent as in the baseline and consent
- Incentivized forecast of air pollution (PM 2.5) levels tomorrow and incentivized guess of the SMS's forecast
- Value elicitation of the SMS forecast service through a bidding game using the BDM method
- Access to information about air pollution and stated satisfaction with the SMS forecast service
- Donation game between EPD and PAQI, and stated preferences for the sources
- Preferences for air quality-related policies via hypothetical scenarios
- Attitudes and behaviors regarding air pollution

- Time use survey and outdoor activities
- Stated mask usage
- Interest in filing complaints about air pollution to government authorities

B.3 Air quality data

Other than the ones discussed in Section 2.1, we collect air quality reading data from the following sources for the forecast model and the intervention.

A government agency called the Urban Unit owns an air quality monitor but has not been consistently publishing readings for the public’s consumption. It is a government-owned yet privately operated entity that addresses urban issues using data in Punjab Province. It was launched as part of a unit in the Planning and Development Department of the provincial government of Punjab in 2005 and was spun off to the private sector with full government ownership in 2012. The unit works on a range of issues pertaining to sustainable urban development, primarily in the realm of environmental services and management. The department owns a high-quality air quality monitor and had previously provided its readings on the banner of their website, but had stopped providing this daily information publicly prior to the beginning of our intervention in early 2023. They have an Environment Dashboard that individuals can sign up for and gain access to historical data on PM2.5 readings, but this data is updated at a lag of 10-15 days. We receive hourly average readings of PM2.5 concentration from the unit’s staff members on a daily basis.

One could also access forecasts based on satellites and meteorological models. One example of such an approach is the Spectral Radiation-Transport Model for Aerosol Species (SPRINTARS), a numerical model that estimates the effect of aerosols on the climatic system via simulations based on an atmosphere-ocean general circulation model called MIROC. The model and estimates have been developed by the Climate Change Science Section at the Research Institute for Applied Mechanics, Kyushu University (Fukuoka, Japan).

SPRINTARS considers both natural and anthropogenic sources of aerosols and categorizes them into suspended particulate matter (SPM), PM2.5, and PM10. Through a collaboration with the model’s developers at Kyushu University, we are able to access the hourly forecasts generated by SPRINTARS. However, we are not aware of any satellite- or other model-based services that actively disseminate air quality information for Lahore or Pakistan.

B.4 Weather Data

We also collect weather data as inputs for the forecast model, as described in further detail in Section 3.2.

- **AccuWeather:** We scrape daily forecasts on maximum and minimum temperatures and precipitation probability from AccuWeather for Lahore at <https://www.accuweather.com/en/pk/lahore/260622/daily-weather-forecast/260622>. AccuWeather uses NOAA’s (National Oceanic and Atmospheric Administration) data and constructs its own forecasts.
- **ASOS:** We also collect detailed meteorological data collected by weather stations at airports. The data sources are called Automated Surface/Weather Observing Systems (ASOS/AWOS) or, more generically, METeorological Aerodrome Reports (METARs). We use a web repository of these data sets hosted by Iowa State University’s Iowa Environmental Mesonet and collect data for a station named “[OPLA] LAHORE(CIV/MIL)” via the following link: https://mesonet.agron.iastate.edu/request/download.phtml?network=PK__ASOS.
- **Weather Underground:** We also collect data on average and minimum atmospheric pressure and daily total precipitation from Weather Underground (URL: <https://www.wunderground.com/weather/pk/lahore>).

C Power Calculations

We estimate the minimum detectable effect sizes on our primary outcomes at 80% probability, with $\alpha = 0.05$. We assume 15 percent attrition on our sample of 1,010. We also make conservative adjustments by dividing the α level by the number of tests for which we are identifying minimum treatment effect sizes.

There are two iterations to our power calculations. First, we identified the number of experimental arms and sample size based on the minimum detectable effect sizes during the design phase in June 2022. Out of the five hypotheses we present in this pre-analysis plan, we had only identified two of them during the design phase (and therefore divide α by 2). We then take sample means and standard deviations from survey data used in Ahmad et al. (2022). The outcomes, sample means, and standard deviations in parentheses are as follows:

1. Willingness-to-pay (WTP) for SMS-based air quality forecasts: 89.6 (45.2)
2. Absolute error of incentivized $t + 1$ forecast of PM2.5 concentration: 43.4 (43.0)

We find that we are able to detect impacts of 0.27 standard deviations, which is equal to PKR 12.3 in the willingness to pay, and $11.7 \mu\text{g}/\text{m}^3$ for PM2.5 concentration.

Second, we re-estimate the minimum detectable effect sizes on the five hypotheses that we pre-specify in this document, using new data from the baseline survey when available. The outcomes, hypotheses, sample means, and standard deviations are:

1. Willingness-to-pay (WTP) for SMS-based air quality forecasts is greater than 0 regardless of the source to which the information is attributed: 89.6 (45.2)

2. Willingness-to-pay (WTP) for SMS-based air quality forecasts is differentially affected by treatment: 89.6 (45.2)
3. Absolute error of incentivized $t + 1$ forecast of PM2.5 concentration, divided by the truth, is differentially affected by treatment: 0.72 (0.42)
4. Perceived accuracy of air-quality information source as the absolute error of incentivized guess of the SMS's forecast is differentially affected by treatment: N/A
5. the amount out of PKR 100 donated to a government agency for an environmental cause, as opposed to the NGO, is differentially affected by treatment: 50.1 (15.0)

For hypotheses 1. and 2., we use the sample statistics from Ahmad et al. (2022) as we do not collect these outcomes in the baseline of this study. We do not have relevant statistics available from either the baseline or from Ahmad et al. (2022) for hypothesis 3., but we expect the outcome variable for it to have a similar distribution to the one for hypothesis 3..

We find that we are able to detect impacts of 0.43 standard deviations, which equals PKR 19.4 in the willingness to pay (for hypothesis 2.), 0.18 for hypothesis 3., and 6.4 for hypothesis 5.. For the test of means for hypothesis 1., we find that we are powered to detect that willingness to pay is greater than PKR 3.6.

Although the minimum detectable impact is fairly large in terms of standard deviations, the treatment effect sizes are relatively small in the outcomes' units. Furthermore, there are several reasons why our assumptions may not hold, or statistical precision could be improved. First, we plan to improve precision by including controls selected via a double-post-selection method using LASSO. Assuming a 30-percent reduction in standard errors, the minimum detectable effects would be 0.30 standard deviations. Second, the willingness-to-pay statistic from Ahmad et al. (2022) may be outdated after two years of high inflation.

D Specification: heterogeneous treatment effects

D.1 Measures of the dimensions of heterogeneity

To measure the dimension of heterogeneity on baseline preferences for, and beliefs about, the sources of air quality information, we use the following proxies:

1. donation share of PKR 100 between government’s environmental agency vs. NGO that tackles air pollution
 - For categorical variables, code as “more to government,” “more to NGO,” and “50-50” or into 10-rupee bins
2. Relative overall approval of government vs. citizen sources: difference in Likert-scale approval measures for the government and NGOs for their air quality information services.
 - For a categorical variable, code as “government-leaning” if the respondents’ Likert-scale approval measure for the government is greater than that for the NGO, “NGO-leaning” if vice versa, and “neutral” if they equally approve the two sources
3. Relative beliefs on the accuracy of government vs. citizen sources: difference in Likert-scale measures for the government and NGOs for their air quality information’s accuracy.
 - For a categorical variable, code as “government-leaning” if the respondents’ Likert-scale approval measure for the government is greater than that for the NGO, “NGO-leaning” if vice versa, and “neutral” if they equally approve the two sources

For robustness, we also consider other definitions of baseline preferences and beliefs, such as the original Likert scales used to construct the proxies above, as well as the respondents’ primary news sources’ political leanings.

For the dimension of heterogeneity on baseline beliefs about air quality and its deviation from the truth, we use the following proxy:

- baseline outcome variable 4.2: absolute error of incentivized $t+1$ forecast of PM2.5 levels.

We also use several other definitions of baseline beliefs to test, for instance, asymmetry based on the direction of the error.

D.2 Estimating equations

The estimating equation to identify the linear ITT effect is as follows:

$$Y_i = \alpha + Z_i\beta + Z_iH_i\theta + H_i\delta + \mathbf{X}_i\boldsymbol{\gamma} + \varepsilon_i \quad (3)$$

H_i is the relevant dimension of heterogeneity as a continuous variable and Z_i the treatment assignment variable that is 1 for the Government arm. We interpret the coefficients $\hat{\beta}$ and $\hat{\theta}$ as estimates of average treatment and heterogeneous treatment effects, respectively.

We also estimate a model where the dimension of heterogeneity is categorical. The estimation equation is as follows:

$$Y_i = \alpha + Z_i\beta + \sum_{j \in J} Z_iH_i\theta_j + \sum_{j \in J} H_i\delta_j + \mathbf{X}_i\boldsymbol{\gamma} + \varepsilon_i \quad (4)$$

H_i is the relevant dimension of heterogeneity as a categorical variable, and each category is denoted as j . We interpret the coefficients $\hat{\beta}$ and $\hat{\theta}_j$ as estimates of the average treatment effect and heterogeneous treatment effect for a group $H_i = j$, respectively.

E Conceptual framework

We specify a consumer’s utility function and how information from a given source alters their beliefs and utility. In our model, consumers hold beliefs about the state variable, i.e., the air quality level. They value the forecast information, with which they can take better mitigation measures against air pollution. We express their value of accessing the forecast information from a given source as a utility function.¹⁶ Consumers also hold beliefs about signal quality (i.e., the accuracy of the SMS forecast) about a source, as well as beliefs and preferences about a source that is not tied to the signal itself. Such beliefs and preferences factor into the utility function as attributes. We exogenously vary the source to which we attribute the signal in our experiment.

E.1 Set-up

The state variable over which consumers form beliefs and receive signals is the air quality for the next day, denoted as q_{t+1} . There are two sources of signals for air quality $s \in \{G, P\}$, government and NGO, respectively. The sources send out SMS forecasts, i.e., signals of air quality for day $t+1$ on day t , denoted as $f_{s,t}^{t+1}$. We model the consumer’s willingness to pay for information from source s when they have received signals up to day t from $a \in \{G, P\}$. a is the information source to which the consumer is already exposed on day t and need not equal s .

On day t , consumer i holds beliefs about the air quality level tomorrow ($E_{i,t}(q_{t+1})$). Each day, they receive a signal on air quality levels tomorrow, $f_{s,t}^{t+1}$, and update their beliefs about the air quality for the next day. Before they receive the signal, they also have beliefs about the signal quality of the SMS forecast for tomorrow ($E_{i,t}(f_{s,t}^{t+1})$). Consumers’ beliefs may also be updated based on the information they have received up to day t (ι) from source a , which we denote as $z(\iota, a)$.

¹⁶For simplicity, we implicitly assume that the cost of accessing any forecast services is sufficiently high that they would only consume information from one source.

E.2 Utility function

A consumer gains utility by accessing an air quality forecast for day $t + 1$ from source $s \in \{G, P\}$ on day t . Consumer i 's utility, $u_{i,a,t}^s$, is defined as follows:

$$u_{i,a,t}^s = \alpha_i + \beta_t + \delta g(E_{i,t}[f_{s,t+1}], E_{i,t}[q_{t+1}]; z(\iota, a)) + \theta b_{i,t}^s(z(\iota, a)) + \epsilon_{i,s,t} \quad (5)$$

In the utility function, the constant individual term is expressed as α_i . The time-varying term β_t captures the value of information that varies by the day and observable conditions they face. The function $g()$ expresses the consumer's beliefs about signal quality $E_{i,t}[f_{s,t+1}]$ conditional on their belief about air quality $E_{i,t}[q_{t+1}]$ and its accuracy. Furthermore, consumers may have a preference for a source s that is not tied to the signal itself, which we express as $b_{i,t}^s(z(\iota, a))$. Lastly, we include an i.i.d. error term $\epsilon_{i,s,t}$.

E.3 Beliefs about signal quality and updating

In Equation 5, $g()$ is an unspecified function. We introduce an additional structure about $g()$ and the belief-updating process to approximate $g()$ as a linear function of the consumer's own forecast accuracy and their belief about the SMS forecast's signal quality. We elicit both of these measures in incentivized games from the surveys, allowing us to map our conceptual framework to the empirical tests.

We first note that the value of the SMS forecast to a consumer should depend on the accuracy of their own beliefs about air quality without access to the forecast. We define this measure to be $E[|q_{t+1} - E_{i,t}(q_{t+1}; z(\iota, a))|]$. For instance, the value of a signal may be greater for individuals with noisier beliefs without access to the signal. In such a case, we should expect:

$$\frac{dg}{dE[|q_{t+1} - E_{i,t}(q_{t+1}; z(\iota, a))|]} > 0 \quad (6)$$


We also assume that the value of the SMS forecast to a consumer should

depend on the additional signal that they believe the forecast provides, conditional on their own beliefs about air quality without the forecast. We define such a measure to be the following:

$$\begin{aligned} E[|q_{t+1} - E_{i,t}(q_{t+1}; f_{s,t+1}, z(\iota, a))|] - E[|q_{t+1} - E_{i,t}(q_{t+1}; z(\iota, a))|] \\ = |E_{i,t}(q_{t+1}; f_{s,t+1}, z(\iota, a)) - E_{i,t}(q_{t+1}; z(\iota, a))| \end{aligned} \quad (7)$$

The consumer engages in a Bayesian belief-updating process about the expected value of q_{t+1} based on the signal, $f_{s,t+1}$. Appendix Figure E.1 describes such a process. Suppose that the SMS forecast signal θ is drawn from a Normal distribution of the true air quality level, whose probability density function is $h(\cdot)$. The priors and posterior beliefs of the distribution are denoted as $\hat{h}(\theta; z(\iota, a))$ and $\hat{h}(\theta; f_{s,t+1}, z(\iota, a))$, respectively. Then the extent to which $\hat{h}(\cdot)$ shifts toward $h(\cdot)$ depends on the deviation of $f_{s,t+1}$ from the mean of the prior distribution, as well as the consumer's perception about the signal quality.

Figure E.1: Prior and posterior beliefs of h based on signal $f_{s,t+1}$



As such, we conjecture that the extent to which the beliefs shift in response to the signal is proportional to the difference between the prior belief about the expected air quality and the SMS signal. In other words, the following is true to an approximation.

$$|E_{i,t}(q_{t+1}; f_{s,t+1}, z(\iota, a)) - E_{i,t}(q_{t+1}; z(\iota, a))| \propto |E_{i,t}(f_{s,t+1}; z(\iota, a)) - E_{i,t}(q_{t+1}; z(\iota, a))| \quad (8)$$

Based on the additional structure introduced above, we can rewrite $g()$ as a linear Taylor approximation in Equation 9. In our field experiment, we measure each of the components of the equation by exogenously varying a .

$$u_{i,a,t}^s = \alpha_i + \beta_t + \gamma * |E_{i,t}(f_{s,t+1}; z(\iota, a)) - E_{i,t}(q_{t+1}; z(\iota, a))| + \omega * E[|q_{t+1} - E_{i,t}(q_{t+1}; z(\iota, a))|] + \theta b_{i,t}^s(z(\iota, a)) + \epsilon_{i,s,t} \quad (9)$$

The utility function allows us to set up a framework for the empirical exercise we conduct. We map the utility function to the hypotheses in the next subsection.

E.4 Pre-specified hypotheses and their links to the utility function

We define five pre-specified hypotheses based on outcomes 4.1 through 4.4, as defined in Section 4. These hypotheses address whether the willingness to pay measure is greater than zero and whether the four primary outcome variables have different levels between the two treatment arms. We organize pre-specified hypotheses and map them to components of the utility function specified as Equation 9.

The left-hand side variable ($u_{i,G,t}^G$ or $u_{i,P,t}^P$) is estimated via the bidding game using the Becker-DeGroot-Marschak (BDM) method. The bid $v_{i,s,t}^a$ in this game would maximize expected utility if $v_{i,s,t}^a = u_{i,a,t}^s$. As such, we observe

distributions of $u_{i,G,t}^G$ and $u_{i,P,t}^P$. We also observe $u_{i,G,t}^P$ and $u_{i,P,t}^G$ through a hypothetical willingness-to-pay survey question in which we asked about the respondent's demand for information coming from a source to which they are not assigned.

We also observe the right-hand-side components of the utility function through survey responses and incentivized elicitation. The consumer's belief about signal quality $E_{i,t}(f_{s,t+1}; z(\iota, a))$ is elicited through the forecast game of SMS forecast and $E_{i,t}(q_{t+1}; z(\iota, a))$ through the forecast game of the air quality level the next day. We observe the realized air quality level of the next day q_{t+1} through air quality readings. We also observe proxies of $b_{i,t}^s(z(\iota, a))$ through stated preferences and beliefs measures about the service quality of each of the sources.¹⁷

The following are the five pre-specified null hypotheses with links to the primary outcome variables and components of the utility function shown as Equation 9:

1. The demand for air quality information is equal to zero regardless of the treatment assignment group. This hypothesis is tested on outcome 4.1 and equates to $u_{i,G,t}^G = 0$, $u_{i,P,t}^P = 0$ in the conceptual framework.
2. The demand for air quality information is not different between the treatment (NGO) and control (government) groups. This hypothesis is tested on outcome 4.1 and equates to $u_{i,G,t}^G = u_{i,P,t}^P$ in the conceptual framework.
3. Treatment does not differentially affect beliefs about air quality relative to control. This hypothesis is tested on outcome 4.2 and equates to $E[|q_{t+1} - E_{i,t}(q_{t+1}; z(\iota, G))|] = E[|q_{t+1} - E_{i,t}(q_{t+1}; z(\iota, P))|]$ in the conceptual framework.

¹⁷We do not observe individual component α_i in the data. As such, we average it out between treatment groups, relying on balance in individual characteristics from randomization. We also control for β_t based on the date of the endline survey.

4. Treatment does not affect the perceived accuracy of the air-quality information source relative to control. This hypothesis is tested on outcome 4.3 and equates to $|E_{i,t}(f_{s,t+1}; z(\iota, G)) - E_{i,t}(q_{t+1}; z(\iota, G))| = |E_{i,t}(f_{s,t+1}; z(\iota, P)) - E_{i,t}(q_{t+1}; z(\iota, P))|$ in the framework.
5. Treatment does not affect relative preferences between information sources. This hypothesis is tested on outcome 4.4 and equates to $b_{i,t}^G(z(\iota, G)) = b_{i,t}^P(z(\iota, G))$ and $b_{i,t}^P(z(\iota, P)) = b_{i,t}^G(z(\iota, P))$ in the conceptual framework.

E.5 Mapping the empirical results to the conceptual framework

Our empirical results show that, although the source does not differentially affect the recipients' demand for air quality information, it affects several underlying beliefs and preferences. To put further structure to our findings, we map our empirical results to Equation 9. We follow the links between pre-specified hypotheses and the conceptual framework highlighted in Section E.4 and identify which attributes of the utility function shift in response to making salient the information source.

We start with measures of the demand for air quality information from a given source. Empirical results in Section 6.3 show that there is a positive demand for air quality information, i.e., $u_{i,G,t}^G > 0$, $u_{i,P,t}^P > 0$. However, we fail to reject the null that there is no differential willingness to pay between treatment groups, i.e., $u_{i,G,t}^G = u_{i,P,t}^P$. Yet, it is unclear if and how the treatment affects each attribute of the utility function and to what extent the changes in beliefs matter to the overall utility. In other words, after averaging out individual fixed effects via randomization and controlling for time fixed effects, failing to reject the null of $u_{i,G,t}^G = u_{i,P,t}^P$ implies the following equation:

$$\begin{aligned}
& \gamma * |E_{i,t}(f_{G,t+1}; z(\iota, G)) - E_{i,t}(q_{t+1}; z(\iota, G))| + \omega * E[|q_{t+1} - E_{i,t}(q_{t+1}; z(\iota, G))|] + \theta b_{i,t}^G(z(\iota, G)) \\
= & \gamma * |E_{i,t}(f_{P,t+1}; z(\iota, P)) - E_{i,t}(q_{t+1}; z(\iota, P))| + \omega * E[|q_{t+1} - E_{i,t}(q_{t+1}; z(\iota, P))|] + \theta b_{i,t}^P(z(\iota, P))
\end{aligned} \tag{10}$$

Equation 10 suggests that, even if the right and left-hand sides are equal, individual components of the equations could be differentially affected by treatment. To highlight this point we estimate correlational relationships between the willingness to pay and other prespecified outcomes by treatment arm, with results shown on Appendix Table E.1. For the NGO arm, we find a negative correlational relationship between the willingness to pay and the SMS error measure, i.e., beliefs about signal quality $|E_{i,t}(f_{P,t+1}; z(\iota, P)) - E_{i,t}(q_{t+1}; z(\iota, P))|$. This relationship, however, is absent for the Government arm. These results are consistent with other empirical results that consumers in the Government arm believe lower signal quality, but do not have lower willingness to pay.

Next, we identify which belief measures in Equation 10 are differentially affected by the information source. Results in Section 6.4 show that there are no differential beliefs about air quality levels between government and NGO sources, i.e., $E[|q_{t+1} - E_{i,t}(q_{t+1}; z(\iota, G))|] = E[|q_{t+1} - E_{i,t}(q_{t+1}; z(\iota, P))|]$. However, results from Section 6.5 show that those assigned to the government arm believe in worse SMS signal quality than those in the NGO arm, i.e., $|E_{i,t}(f_{G,t+1}; z(\iota, G)) - E_{i,t}(q_{t+1}; z(\iota, G))| > |E_{i,t}(f_{P,t+1}; z(\iota, P)) - E_{i,t}(q_{t+1}; z(\iota, P))|$.

The results above highlight two possibilities. One is that the effect of differential beliefs about signal quality between two sources is offset by other beliefs, i.e., $\theta b_{i,t}^G(z(\iota, G)) > \theta b_{i,t}^P(z(\iota, P))$, assuming $\gamma < 0$. We do not find evidence of such an offsetting mechanism. Based on stated-preference measures in Tables 3 and 4, we find that the respondents improve their approval of the assigned source in the equal magnitude between treatment arms. Thus, we conclude that another possibility is more likely: respondents do not value the precision of SMS forecasts at the current margin of error.

Lastly, we find that the exogenous attribution to a source increases demand for it relative to other alternative sources driven by improved beliefs about signal quality and other aspects of the information source. Table 2 shows that when we compare within individuals between their assigned source and a hypothetical counterpart, recipients have significantly higher demand for the assigned one, i.e., $u_{i,G,t}^G > u_{i,G,t}^P$ and $u_{i,P,t}^P > u_{i,P,t}^G$. These results imply the following condition for those assigned to the government arm (and symmetrically for those in the NGO one):

$$\begin{aligned} & \gamma * |E_{i,t}(f_{G,t+1}; z(\iota, G)) - E_{i,t}(q_{t+1}; z(\iota, G))| + \theta b_{i,t}^G(z(\iota, G)) \\ & > \gamma * |E_{i,t}(f_{P,t+1}; z(\iota, G)) - E_{i,t}(q_{t+1}; z(\iota, G))| + \theta b_{i,t}^P(z(\iota, G)) \end{aligned} \quad (11)$$

Condition 11 shows that the within-person difference in the valuation of government v.s. NGO sources may come from their belief in signal quality or in other unrelated factors. Unfortunately, we cannot empirically observe $E_{i,t}(f_{P,t+1}; z(\iota, G))$. However, Condition 11 is consistent with the results from Tables 3 and 4, which show that respondents increase their approval of their assigned source in terms of its signal quality, reliability (e.g., promptness of the daily forecast SMS messages), and on their overall satisfaction in similar magnitudes.

Furthermore, we correlationally demonstrate which aspects of Condition 11 lead the inequality, particularly their valuations of signal quality vis-à-vis reliability and other attributes. For respondents assigned to the Government arm, we observe $u_{i,G,t}^G$ and $u_{i,G,t}^P$ based on incentivized and hypothetical willingness-to-pay measures. $u_{i,G,t}^G - u_{i,G,t}^P > 0$ underpins the inequality in Condition 11. We regress $u_{i,G,t}^G - u_{i,G,t}^P$ on the differences in Likert-scale approval of government source to the NGO one in terms of signal quality, reliability, and overall service quality. We also have equivalent measures for those in the NGO arm, i.e., $u_{i,P,t}^P$, $u_{i,P,t}^G$, and the corresponding Likert-scale measures.

Appendix Table E.2 shows the results. We find that, for both Government and NGO arms, the Likert-based measure of reliability is more strongly

correlated with the difference in willingness to pay than accuracy or overall satisfaction. These correlational results indicate that consumers update their beliefs about the reliability and timeliness of the assigned source more strongly than about other aspects of the service, leading to an increased demand for the said source.

Table E.1: Relationships between prespecified outcome variables

	(1)	(2)
	WTP: Government	WTP: NGO
Forecast error	6.93*	5.89
	(4.18)	(5.16)
SMS error	0.080	-0.39**
	(0.084)	(0.18)
Donation gov't	-0.24	-0.45
	(0.26)	(0.29)
Observations	494	494

Notes: Model: PDSLASSO. The table shows results of regressing the endline willingness to pay for two months of SMS air quality forecasts on forecast error, SMS error, and share of donations to the government. Column 1 is restricted to those in the Government arm, and Column 2 those in the NGO arm. “Forecast error”: the absolute difference between their forecast air pollution level on the next day and the actual reading, divided by the actual reading. “SMS error”: the absolute difference between their forecast air pollution level and their guess of the SMS forecast on the next day. “Donation gov’t”: amount out of PKR 100 donated to the government source. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.

Table E.2: Relationships between differences between the sources in willingness to pay and beliefs about service quality

	(1)	(2)
	diff(WTP): Government	dif(WTP): NGO
Diff(Govt - NGO Reliable)	1.97** (0.78)	-4.91** (2.06)
Diff(Govt - NGO Accurate)	1.56 (1.09)	-0.64 (0.93)
Diff(Govt - NGO Approve)	-1.76 (1.07)	1.62 (1.28)
Observations	494	492

Notes: Model: PDSLASSO. The table shows results of regressing the differences between sources on the willingness to pay and the differences between sources on beliefs about service quality. Column 1 is restricted to those in the Government arm, and Column 2 those in the NGO arm. “diff(WTP): Government”: differences between the elicited willingness to pay for the forecast service that is experimentally associated with the Government and a hypothetical willingness to pay if the information came from the NGO. “diff(WTP): NGO”: differences between the elicited willingness to pay for the forecast service that is experimentally associated with the NGO and a hypothetical willingness to pay if the information came from the government. “Diff(Govt - NGO Reliable)”: differences in stated beliefs (measured in the Likert scale) that the Government source is reliable and on time, relative to the NGO one. “Diff(Govt - NGO Accurate)”: differences in stated beliefs (measured in the Likert scale) that the Government source is accurate, relative to the NGO one. “Diff(Govt - NGO Approve)”: differences in overall approval (measured in the Likert scale) of the Government source, relative to the NGO one. Standard errors are reported in parentheses. All regressions include randomization-strata fixed effects, and heteroskedasticity-robust standard errors are used. Two-tailed significance: $p < 0.1^*$; $p < 0.05^{**}$; $p < 0.01^{***}$.