

NBER WORKING PAPER SERIES

MISPERCEPTIONS ABOUT AIR POLLUTION:
IMPLICATIONS FOR WILLINGNESS TO PAY AND ENVIRONMENTAL INEQUALITY

Matthew A. Tarduno
Reed Walker

Working Paper 34116
<http://www.nber.org/papers/w34116>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2025, Revised June 2026

We thank Michael Anderson, Josh Apte, Patrick Baylis, Catie Hausman, Kelsey Jack, Ricardo Perez-Truglia, David Schonholzer, Katherine Wagner, and seminar participants at the University of British Columbia, the Harvard University Center for the Environment, University of Pittsburgh, and the University of Illinois Urbana-Champaign for useful comments and discussions. Funding for the pilot for this research from Berkeley's Xlab is gratefully acknowledged. Suvy Qin and Ray Kim provided excellent research assistance for this project. We also thank Sara Johns, Jesse Buchsbaum, Matthew Suandi, and Livia Alfonsi for helpful feedback on the survey design, and Kristie H. Dinh for diligent assistance in distributing the survey rewards. The study design was pre-registered on the AEA repository under AEARCTR-0010435 and AEARCTR-0011970. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Matthew A. Tarduno and Reed Walker. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Misperceptions About Air Pollution: Implications for Willingness to Pay and Environmental Inequality

Matthew A. Tarduno and Reed Walker

NBER Working Paper No. 34116

August 2025, Revised June 2026

JEL No. H4, Q5

ABSTRACT

This paper explores whether misperceptions about air pollution contribute to environmental inequality in the United States. We use a two-part survey experiment to elicit respondents' beliefs about local air quality and pollution's effects on life expectancy. We document how misperception differs across demographic groups and then how this misperception affects willingness to pay (WTP) for cleaner air. Since misperception or beliefs may be correlated with other unobservable determinants of WTP, we randomly show selected participants customized information about their actual air pollution. This allows us to trace out how experimentally induced changes in beliefs affect WTP for air quality. Our results suggest significant misperceptions about air pollution in the US. Respondents, on average, overestimate both their air pollution exposure and its impact on life expectancy. Beliefs about relative air pollution are not systematically biased but are noisy. Despite some differences in misperceptions between Black and White respondents, counterfactual exercises do not suggest that rectifying these misperceptions would meaningfully close the observed gap in WTP and/or pollution exposure.

Matthew A. Tarduno
University of Illinois Chicago
tarduno@uic.edu

Reed Walker
University of California, Berkeley
Haas School of Business
Department of Economics
and NBER
rwalker@berkeley.edu

A randomized controlled trials registry entry is available at
<https://www.socialscienceregistry.org/trials/10435> and
<https://www.socialscienceregistry.org/trials/11970>

1 Introduction

Environmental inequality—the disproportionate exposure of low-income households and racial minorities to pollution—is one of the most persistent features of the American environmental landscape. Landmark studies in the 1980s and 1990s documented these disparities (see e.g., [Bullard, 1993](#)), and despite fifty years of progress on air and water quality, the gradient between economic prosperity and pollution exposure remains. Understanding why environmental inequality persists is a central question in environmental economics and a prerequisite for effective policy.

The literature has advanced several hypotheses. Differences in income may lead to differential sorting across neighborhoods of varying environmental quality ([Depro, Timmins, and O’neil, 2015](#); [Colmer, Qin, Voorheis, and Walker, 2026](#)). Heterogeneous preferences or willingness to pay for amenities may play a role ([Kuminoff, Smith, and Timmins, 2013](#); [Banzhaf, Ma, and Timmins, 2019](#); [Been and Gupta, 1997](#)). Discriminatory housing practices ([Christensen and Timmins, 2022](#)) and pollution siting decisions ([Pastor, Sadd, and Hipp, 2001](#); [Wolverton, 2009](#)) may directly generate or reinforce disparities. Each of these channels has received substantial empirical attention.

A hypothesis that has received far less empirical scrutiny is that *differences in information*—in what people know about their pollution exposure and its consequences for health—may contribute to observed disparities ([Hausman and Stolper, 2021](#)). The logic is straightforward: if households are unaware that they face disproportionate pollution, or if they underestimate the health consequences, they will underinvest in avoidance, exert less political pressure for cleanup, and sort less aggressively away from polluted areas. Information frictions could thus sustain environmental inequality even absent any differences in underlying preferences. Despite its theoretical appeal, this hypothesis has been difficult to test because it requires data that, until now, have not existed: measures of what people actually *believe* about their pollution exposure, how those beliefs compare to reality, and whether beliefs differ across the demographic groups that figure centrally in environmental justice debates.

Beyond environmental inequality, misperceptions about pollution carry implications for how economists measure the demand for environmental quality. A large literature estimates willingness to pay for clean air from housing markets and defensive investments, typically under the assumption that agents are well-informed about local amenities ([Gao, Song, and Timmins, 2021](#); [Kuminoff, Smith, and Timmins, 2013](#)). If households systematically misperceive pollution—overestimating or underestimating exposure or health impacts—then revealed-preference estimates will reflect *perceived* rather than true valuations. The direction and magnitude of this bias depend on the structure of misperceptions, which is ultimately an empirical question.

This paper provides the first comprehensive empirical test of these hypotheses in the United States. We design and implement a large-scale survey experiment that (i) measures beliefs about air pollution along multiple dimensions, (ii) compares the accuracy of those beliefs across racial and income groups, (iii) exogenously shifts beliefs through randomized information provision, and (iv) traces the causal chain from beliefs to willingness to pay for clean air.¹ We elicit three types of beliefs from a nationally representative sample of over 5,000

¹We pre-registered this study with the AEA RCT registry (*AEARCTR-0010435*) and followed a split-sample approach to multiple hypothesis testing described by [Anderson and Magruder \(2017\)](#). This approach reduces the likelihood of false discoveries in settings like ours that involve many hypotheses, but it also allows the researchers some flexibility to perform additional ex-post tests, albeit with lower power. Throughout our paper we describe how we design our survey and analysis with regard to this

respondents: beliefs about *relative* pollution exposure (how does your neighborhood compare to others in your state?), beliefs about pollution *levels* (what is PM_{2.5} concentration in your ZIP code?), and beliefs about the *health impacts* of pollution (how much does air pollution reduce life expectancy?). After recording initial beliefs, we randomly provide respondents with accurate, individualized information about one of these dimensions and measure how beliefs and willingness to pay respond.

Our empirical design allows us to test several specific hypotheses from the environmental justice and environmental valuation literatures. First, do misperceptions about pollution exist, and are they large enough to plausibly affect behavior? Second, do misperceptions differ systematically across racial and income groups in ways that could sustain environmental inequality? Third, does correcting misperceptions causally affect willingness to pay for clean air? And fourth, would moving households to a fully informed state meaningfully close observed gaps in willingness to pay or pollution exposure between demographic groups? In doing so, we are able to recover estimates of willingness to pay for air quality, separately by race and income — decomposing observed gaps in WTP into a component attributable to differential information and underlying preference heterogeneity.

We find that misperceptions are substantial and widespread. Americans overestimate both local air pollution levels and the health impacts of pollution exposure by roughly an order of magnitude. Beliefs about relative exposure—whether one lives in a clean or dirty area—are unbiased on average but exhibit significant compression: residents of the cleanest areas overestimate their pollution, while residents of the dirtiest areas underestimate theirs. Importantly, these misperceptions are *similar* across demographic groups. The one exception is beliefs about pollution exposure between White and Black respondents: White respondents hold more accurate beliefs than Black respondents about both their relative exposure and their exposure in levels, even after conditioning on actual pollution levels. We find no significant differences by race in beliefs about the health impacts of pollution, and no significant differences by income in any of our three belief measures.

Using a real-stakes willingness-to-pay exercise, we estimate that respondents are willing to pay roughly \$1 for a 1% reduction in air pollution at home, or about \$8 per $\mu\text{g}/\text{m}^3$. MWTP is significantly higher among White and high-income respondents, both unconditionally and conditional on observables. We then use our experimental variation to trace the causal effect of beliefs on MWTP. We find that beliefs about the health impacts of pollution are the primary driver: a one-year increase in perceived life-expectancy impact raises MWTP by roughly 28%. Beliefs about pollution exposure—whether measured in levels or relative terms—have no significant effect on MWTP, even though our information treatments successfully shift exposure beliefs. Since most respondents overestimate health impacts, correcting these beliefs *decreases* MWTP—a finding that contrasts with the existing literature, which has largely been conducted in low- and middle-income countries where people tend to underestimate pollution.

What do these findings imply for environmental inequality? Our counterfactual exercises suggest that moving all respondents to fully informed beliefs would reduce average MWTP by nearly half, but would not meaningfully close the Black-White gap in willingness to pay or pollution exposure. Because beliefs are similar across groups, the information channel cannot explain the observed WTP gap—the gap that remains after correcting misperceptions is, if anything, slightly larger in proportional terms. This finding has two implications. First, it

multiple hypothesis protocol.

provides direct evidence against the hypothesis that information frictions are a primary driver of environmental inequality in the United States, at least through the belief channels we measure. Second, and perhaps more constructively, it implies that the demographic gaps in WTP we estimate—which persist after experimentally removing information differences—are more likely to reflect genuine preference heterogeneity (broadly defined to include income effects) than differential awareness. These “informed” WTP estimates, which account for the role of misperceptions, are to our knowledge the first of their kind for air quality in the US, and provide a new empirical input for models of environmental sorting and policy design that require preference parameters by demographic group.

This paper makes several contributions. First, we provide the first direct test of the information hypothesis for environmental inequality, using measures of actual beliefs rather than proxies. We find that while misperceptions are large, they do not differ enough across groups to drive observed disparities—redirecting attention toward income, housing markets, and institutions as the more likely sustaining forces. Second, by experimentally correcting beliefs and measuring the residual WTP gap, we provide among the first credible estimates of preference heterogeneity for environmental quality across race and income groups. The existing literature has lacked the ability to separate preferences from information in revealed WTP; our design addresses this directly. Third, by separately eliciting beliefs about exposure, relative exposure, and health impacts, we identify which dimensions of understanding matter for demand. The pollution-health channel dominates: beliefs about how harmful pollution is matter far more for willingness to pay than beliefs about how much pollution one faces. This has implications for both information policy and the interpretation of existing revealed-preference estimates. Fourth, our MWTP estimates provide a useful US benchmark for the growing literature on demand for clean air.

The rest of this paper proceeds as follows. Section 2 introduces a simple conceptual framework used to highlight the channels by which information provision may affect willingness to pay, and how differences in misperceptions and/or preferences across demographic groups may contribute to observed sorting patterns. Section 3 describes the survey, and Section 4 describes the sample. In Section 5 we describe and discuss the results of the initial survey, the follow-up survey, and the willingness to pay exercise. Section 6 concludes.

2 Conceptual Framework

We develop a simple framework linking beliefs to willingness to pay for cleaner air to clarify how misperceptions and information provision may affect demand for air quality.

Setup

Consider an individual i who derives utility from consumption c and health h , with utility function $U(c_i, h_i)$ that is increasing and concave in both arguments (following Grossman, 1972). Health depends on pollution exposure according to:

$$h_i = \bar{h} - \gamma \cdot e_i \tag{1}$$

where $\bar{h}$ is baseline health, e_i is true pollution exposure, and $\gamma > 0$ is the true health impact per unit of pollution (e.g., life-years lost per $\mu\text{g}/\text{m}^3$).

Individuals do not observe γ or e_i perfectly. Instead, they form beliefs $\tilde{\gamma}_i$ and $\tilde{e}_i$ about the health impact and their exposure, respectively. The individual’s *perceived* health is therefore:

$$\tilde{h}_i = \bar{h} - \tilde{\gamma}_i \cdot \tilde{e}_i \quad (2)$$

We assume individuals can invest in pollution avoidance (e.g., air purifiers, residential location choice) by purchasing exposure reduction Δ at cost p per unit. Substituting perceived health into the utility function, the individual chooses Δ to maximize:

$$U(y_i - p \cdot \Delta, \bar{h} - \tilde{\gamma}_i \cdot (\tilde{e}_i - \Delta)) \quad (3)$$

Demand for Clean Air

The first-order condition yields the individual’s marginal willingness to pay for a unit reduction in pollution exposure:

$$MWTP_i = \frac{U_h(c_i, \tilde{h}_i)}{U_c(c_i, \tilde{h}_i)} \cdot \tilde{\gamma}_i \quad (4)$$

or, normalizing the marginal utility of consumption to one:

$$MWTP_i = U_h(c_i, \tilde{h}_i) \cdot \tilde{\gamma}_i \quad (5)$$

This expression highlights two key points about how beliefs enter demand for clean air.

First, perceived health impacts ($\tilde{\gamma}_i$) and perceived exposure ($\tilde{e}_i$) affect MWTP through distinct channels. The perceived health impact $\tilde{\gamma}_i$ enters directly: it is the perceived health gain from a one-unit reduction in exposure, arising from the chain rule $\partial \tilde{h}_i / \partial \Delta = \tilde{\gamma}_i$. An individual who believes pollution is very harmful will have high MWTP, all else equal. Perceived exposure $\tilde{e}_i$, by contrast, affects MWTP only indirectly—through the marginal utility of health, U_h . If U is concave in h , an individual who perceives high exposure (and therefore low health) will have a higher marginal utility of health, and thus a higher MWTP. However, if utility is approximately linear in health over the relevant range, then U_h is roughly constant and exposure beliefs will have limited effects on MWTP.²

Second, if individuals hold upward-biased beliefs ($\tilde{\gamma}_i > \gamma$ or $\tilde{e}_i > e_i$), then their observed MWTP will exceed their “fully-informed” MWTP. Conversely, if beliefs are downward-biased, observed MWTP will understate true demand. This framework thus provides a mapping between the direction of misperceptions and the direction of bias in revealed-preference estimates of environmental valuations.

Implications for Environmental Inequality

The framework also clarifies how misperceptions could contribute to environmental inequality. Consider two groups, A and B , where group B faces higher true pollution exposure ($e_B > e_A$). Environmental inequality could arise or persist through several belief-related channels:

²In the limit where U is linear in h , U_h is constant and MWTP depends only on $\tilde{\gamma}_i$. This corresponds to the case where the dose-response function is effectively linear from the individual’s perspective.

1. **Differential awareness:** If group B underestimates their relative exposure ($\tilde{e}_B - e_B < \tilde{e}_A - e_A$), they may underinvest in avoidance relative to group A , perpetuating exposure gaps.
2. **Differential health impact beliefs:** If group B has lower perceived health impacts ($\tilde{\gamma}_B < \tilde{\gamma}_A$), they will have lower MWTP even at the same perceived exposure level.
3. **Differential updating:** Even if initial beliefs are similar, groups may respond differently to information due to differences in trust, prior uncertainty, or cognitive resources.

Our empirical design allows us to test each of these channels by (i) measuring beliefs about exposure and health impacts separately, (ii) comparing belief accuracy across demographic groups, and (iii) estimating heterogeneous responses to information provision.

3 Survey Design

We designed an online survey to measure beliefs about air pollution, separately by race, and to test whether willingness to pay for clean air changes when people update their beliefs. The survey had two parts: an initial and a follow-up survey. The initial survey gathered information about respondents’ beliefs about air pollution. We then provided some respondents with customized information about either their air pollution exposure or the health impacts of air pollution. The follow-up survey measured whether individuals updated their beliefs, and whether any changes in beliefs induced by the information provision impacted willingness to pay for clean air. We also conducted a second follow-up to study whether belief updating persists past the first follow-up survey. The following subsection provides an overview of this survey. [Appendix C](#) contains a vignette of the survey instruments, as well as links to copies of the Qualtrics surveys.

3.1 Initial Survey

In the initial survey, we elicited people’s beliefs about local air pollution. We focused on three kinds of beliefs: beliefs about *relative* air pollution exposure, beliefs about air pollution exposure in *levels*, and beliefs about the connection between air pollution and health. We then randomly provided three-quarters of respondents with customized information regarding one of these three questions; the remaining one-quarter of respondents received no information. [Figure A1](#) shows a high-level overview of this survey instrument.

Introduction questions and incentive compatibility

Before eliciting beliefs about air pollution exposure and the health consequences of pollution, we first asked respondents where they live (their home state and ZIP code). These responses were later used to tailor information provision for respondents who were randomized into the treatment arm. We also informed respondents that they would be rewarded for answering questions accurately — the 10% of respondents with answers closest to the correct answers for their ZIP code received a \$2 bonus. These rewards were intended to make our initial elicitation incentive-compatible ([Haaland, Roth, and Wohlfart, 2023](#)).

Eliciting prior beliefs

We asked respondents to share their beliefs about: (Q1) their relative pollution exposure, i.e., how ambient concentrations of $\text{PM}_{2.5}$ in their ZIP code compare to other ZIP codes, (Q2) the level of $\text{PM}_{2.5}$ in their ZIP code (with the Clean Air Act National Ambient Air Quality Standards (NAAQS) for $\text{PM}_{2.5}$ provided as a reference point), and (Q3) the extent to which air pollution affects average life expectancies.

To elicit beliefs about relative pollution exposure, we asked each respondent how they would rate $\text{PM}_{2.5}$ concentrations in their ZIP code relative to other ZIP codes within their state. Respondents answered using a slider that ranged from “Better than 0% of ZIP codes” to “Better than 100% of ZIP codes.” Respondents were required to move the slider to answer the question; there was no default answer. Below the question, we also linked to a ZIP code map of the US, which respondents could click on if they wanted more information about the spatial composition of ZIP codes in their state.

The second elicitation question asked respondents to share their beliefs about their air pollution exposure in levels. We provided a reference point for respondents, because it is unlikely that the average respondent knows how scientists and policymakers measure air quality. The question text read: *The Clean Air Act regulates air quality in the United States. The current national air quality standard for particulate air pollution is 12 micrograms per cubic meter ($\mu\text{g}/\text{m}^3$). Areas with annual average air quality that exceed this threshold are subject to stringent pollution requirements. What do you think the average air pollution level in your ZIP code is?* Respondents reported their answer in a small text box labeled $\mu\text{g}/\text{m}^3$. There were no other instructions, but if respondents provided an answer less than 0 or greater than $300 \mu\text{g}/\text{m}^3$, they were prompted to provide a response between these two values.

Lastly, we elicited beliefs about the relationship between air pollution and health. We first asked respondents a simple yes/no question about whether or not they believe air pollution impacts health: *If average air pollution in the US were 50% higher than it is right now, do you think this would impact US life expectancies?* Respondents who answered “no” were coded as having a prior belief that there is no pollution-life expectancy gradient. Respondents who answered “yes” were then asked an additional question to elicit their prior about the size of this impact. The opening text of this question read: *Smoking, alcohol consumption, and motor vehicle accidents all decrease life expectancy. If average air pollution in the US were 50% higher than it is right now, by how much do you think this would impact life expectancy?* Respondents answered using a slider that ranged from 0 to 4 years, with the reference mortality causes — smoking (2.4 years), alcohol consumption (0.5 years), and vehicle accidents (0.3 years) — labeled along the slider. For this question, we take the “correct” answer to be the mean of two academic studies on the pollution-mortality gradient, [Anderson \(2020\)](#) and [Krewski, Jerrett, Burnett, Ma, Hughes, Shi, Turner, Pope III, Thurston, Calle et al. \(2009\)](#). These studies suggest that if average air pollution in the US were 50% higher than it is right now, life expectancies would decrease by 0.09 and 0.27 years, respectively.

Information Provision

The primary goals of our survey were to (a) test whether individuals misperceive their air pollution exposure or its health impacts, and (b) test whether correcting misperceptions changes willingness to pay for clean air. To study the latter, we randomly provided information to a set of participants in order to isolate exogenous

variation in follow-up beliefs about air pollution. The second stage relates these exogenous differences in beliefs to differences in willingness to pay for clean air.

After eliciting priors, we provided a random subset (75%) of participants with information about *one* of the above questions (Q1 through Q3).³ This information was source-randomized: conditional on the type of information a respondent was randomized to receive (e.g., information about the average particulate matter air pollution level in their ZIP code), this information came from one of two possible sources. Each piece of information was based on academic research, and a link to the academic study was included in the information provision screen. To the extent that individuals update their beliefs in response to information, our survey design yields several instruments for follow-up beliefs. We investigate two such instruments in this paper.

First, we instrument for follow-up beliefs using a treatment indicator interacted with the initial belief gap: $1[treated] \times (signal_i - prior_i)$. This design does not explicitly take advantage of the source-randomized information; it instead isolates variation in posteriors generated by initial misperceptions interacted with whether an individual randomly received the treatment.

We also investigate whether source-randomization induces variation in follow-up beliefs. If participants' beliefs are very responsive to information, or information sources differ significantly, source-randomization allows the researcher to use the information *source* as an instrument for posterior beliefs (as in Bottan and Perez-Truglia, 2022).⁴ Specifically, when regressing a measure of WTP on beliefs, one can instrument for beliefs using $1[treated] \times (signal_i - alt_i)$, where alt_i is the alternative signal that a respondent would have been shown in the alternative realization of the source-randomization.

Although source-randomization is designed to address concerns about spurious updating in information provision, it also carries implications for the interpretation of the local average treatment effect: In designs that use a single information signal rather than a source-randomized signal, the IV “compliers” are those who initially held incorrect beliefs. Alternatively, in a source-randomized setup, compliers may initially hold beliefs that are roughly correct, but they may be induced to update depending on what piece of information they are shown. These instruments therefore carry subtly different interpretations: source-randomization can be taken as identifying the impact of changing beliefs; whereas simple information treatments can be taken as identifying the impact of correcting misperceptions.⁵

To foreshadow our empirical approach below, source-randomization generates relatively little variation in posterior beliefs, whereas the $1[treated] \times (signal_i - prior_i)$ strongly predicts posterior beliefs in two of our three outcomes. As such, we use the simpler “signal - prior” instrument in our main instrumental variable estimates of willingness to pay for clean air.

³Tables A1 and A2 show balance across treatment arms by race and income, respectively. Our treatment is balanced by income, and balanced for race groups with large numbers of respondents: White, Black or African, Asian, and 2 or more races.

⁴Source-randomization is a non-deceptive approach to inducing variation in beliefs among those who receive information. An example of a deceptive information provision survey is one that adds noise to the information provided to respondents, and then uses only the noise as an instrument for follow-up beliefs in subsequent regressions (Haaland, Roth, and Wohlfart, 2023; Bottan and Perez-Truglia, 2022).

⁵In some contexts the choice of instrument does not appear to matter, as in Bottan and Perez-Truglia (2025). In theory, however, these two approaches could yield different results for the reasons described.

Information Sources

We provided information about PM_{2.5} pollution exposure percentile (Q1) and levels (Q2) using two separate satellite-based measures of US particulate matter exposure: Di, Kloog, Koutrakis, Lyapustin, Wang, and Schwartz (2016), and Van Donkelaar, Martin, Li, and Burnett (2019). These two data sources are plotted in Figures A2 and A3. Both of these papers combine air pollution monitoring data with measures of aerosol optical depth (AOD) derived from satellite imagery to estimate particulate matter concentrations. Van Donkelaar et al. (2019) use “a chemical transport model (GEOS-Chem) and satellite observations of aerosol optical depth, and statistically fuse these estimates with ground-based observations using a geographically weighted regression over North America.” Di et al. (2016) use monitoring data and AOD as inputs to a neural network, testing the model with hold out monitors. Both papers provide a raster of PM_{2.5} at 1 km × 1 km grid cells⁶ (the coverage for Van Donkelaar et al. (2019) is North America; the coverage for Di et al. (2016) is the continental United States). For Di et al., we take the average of pollution levels from 2013 to 2015; for Van Donkelaar et al, we take the average of 2016 through 2018.⁷

To provide source-randomized information about life expectancy, we used pollution-mortality concentration response functions from Anderson (2020) and Krewski et al. (2009). These studies estimate that a 10 $\mu\text{g}/\text{m}^3$ increase in PM_{2.5} exposure is associated with a decrease in life expectancy of 0.25 and 0.77 years, respectively. Using a US population-weighted average PM_{2.5} exposure of 7.1 $\mu\text{g}/\text{m}^3$ from the 2022 Air Quality Life Index (Greenstone and Fan, 2018), these estimates suggest that a 50% increase in average US PM_{2.5} would decrease life expectancy by 0.09 or 0.27 years, respectively. We deliberately structured our survey question to concern the impacts of pollution on mortality *on average* to avoid conflating misperception with private information (e.g., health conditions that may lead individuals to have idiosyncratic risks related to air pollution exposure).

For each possible treatment, we followed the same format for providing information. We first repeated each respondent’s answer, we then described the estimate from the research. For example, if a respondent received information about the first question, the information screen would read: *You answered that pollution in your ZIP code was better than [Q1 answer]% of ZIP codes within your state. According to satellite data, the pollution level in your ZIP code is better than [source-randomized information]% of ZIP codes within your state.*

To increase the likelihood that respondents read and remembered this information, we included a timer on this screen so that the button to continue the survey did not appear until 4 seconds had passed. We also included a link which, when clicked, opened a new tab that (a) briefly described the research behind the given piece of information, and (b) provided another link to the actual publication.

⁶Although satellite-based measures of air quality will inevitably suffer from some prediction error (Fowlie, Rubin, and Walker, 2019), both of these data sources provide a significant improvement in spatial coverage over the EPA’s air monitor network. Information provision using monitor-based measurements alone would require either (a) imputing air pollution from far-away measurements, or (b) restricting the sample to areas with monitors, which would likely lead to an unrepresentative sample of respondents by excluding areas that do not have access to high-quality monitoring data.

⁷Although our data are not contemporaneous with our survey (winter and spring of 2023), note that 3-year ZIP Code-level air pollution measures are highly correlated over time. For example, using data from Di et al. (2016), the ZIP code-level correlation coefficient between the 2004-2006 average PM_{2.5} and the 2014-2016 average PM_{2.5} is 0.95.

3.2 Follow-Up Survey

Two weeks after completing the initial survey, respondents were eligible to take the follow-up survey. The purpose of this follow-up survey was to determine whether the information provision at the end of the initial survey induced changes in beliefs, and if so, to test whether changes in beliefs led to changes in willingness to pay for clean air. Figure A4 provides a high-level overview of this survey instrument.

Eliciting Posterior Beliefs

The first section of the follow-up survey repeated Q1 through Q3 from the initial survey for both respondents that received information as well as for respondents in the control group. Responses to these questions allow us to understand belief updating, and provide the first stage for our IV estimates of the impacts of beliefs on MWTP.

Eliciting WTP for Reductions in Air Pollution

The final section of the follow-up survey was designed to allow us to estimate each respondent’s marginal willingness to pay for air quality improvements: We asked each respondent to complete two choice exercises involving air purifiers that differed only in their cleaning capacity. By comparing willingness to pay across these two purifiers, we can estimate each individual’s marginal willingness to pay for pollution reductions—that is, how much each respondent is willing to pay for a 1% reduction in air pollution at home. Importantly, this approach allows us to disentangle willingness to pay for improvements in air quality from willingness to pay for an air purifier itself. We see this as an improvement on methods used elsewhere in the literature, which often use demand for goods like air purifiers or masks as a proxy for clean air, or attempts to back out willingness to pay for clean air using the demand for a single good.

More specifically, this exercise takes advantage of an attribute of air purifiers known as the *clean air delivery rate* (CADR). An air purifier’s effectiveness at reducing indoor pollution levels depends on (a) the air purifier’s clean air delivery rate, (b) the volume of the house, and (c) the number of indoor-outdoor air changes per hour that results from structural leaks in the house. For “leaky” spaces, the indoor-outdoor air turns over at a rate of 2 air changes per hour; for “tight” spaces, the indoor-outdoor air turns over at a rate of 0.1 air changes per hour.

These choices were presented as follows: For each individual, we first asked them how large their home or apartment is, in square feet. We converted this measure to volume assuming 9-foot ceilings. We then showed each respondent two air purifiers: *Purifier A*, which has a CADR of 70 m³/hr, and *Purifier B*, which has a CADR of 240 m³/hr. We did not directly show the respondents the CADR, or information about the brand of the air purifier. Instead, we used the information about the respondent’s home size and the CADR of the air purifier to provide individualized information about the effectiveness of each of these air purifiers.⁸ Figure 1 shows an example of this customized information.

⁸We assume an indoor-outdoor air turnover rate of one change per hour for all homes.

Figure 1: Example of Customized Air Purifier Information

*An air purifier’s effectiveness depends on the **capacity of the air purifier** and the **size of the home** that it is in. In the following questions, you will be asked to make decisions about these two types of air purifiers:*

Air Purifier A

*Based on your home size, this air purifier would reduce indoor air pollution in your home by **23%***



Air Purifier B

*Based on your home size, this air purifier would reduce indoor air pollution in your home by **51%***



Notes: This is an example of the information that a respondent could have been shown about the effectiveness of the two air purifiers that we would offer them in an incentivized price list experiment. We use each air purifier’s clean air delivery rate plus the respondent’s reported home size to provide these estimates of air purifier effectiveness; this particular example reflects information shown to an individual with a 900 square foot apartment.

For each air purifier, we elicited respondents’ willingness to pay using a multiple price list (MPL). A single question on one of these price lists asked respondents to choose between an air purifier and a gift card of a given value. To incentivize respondents to answer the MPL questions according to their true preferences, we notified respondents that 1 in 50 participants would randomly receive one of the items that they selected. Figure A5 provides an example of this multiple price list elicitation.

We followed a procedure proposed by [Jack, McDermott, and Sautmann \(2022\)](#) in order to address inconsistencies often found in survey responses using multiple price lists (e.g., preference reversal). This procedure includes randomizing the order of prices (ascending or descending) and the order of items across individuals (i.e., air purifier in the first column or gift cards in the first column). We also model respondents’ behavior with error to address issues that emerge from “multiple-switching behavior.”⁹ If all respondents switched choices exactly once, then WTP would lie on the interval between the gift card values where they switched. We follow [Jack, McDermott, and Sautmann \(2022\)](#) by pooling observations and estimating a discrete choice model, which allows researchers to use data from others’ choices to estimate WTP for individuals who switch more than once.

In practice, we estimate WTP using a logit regression where: (1) each individual choice is modeled as a separate event, (2) fixed effects for each randomization arm control for which item was shown first and whether prices were ascending or descending, and (3) the sum of the intercept and an individual’s fixed effect corresponds to an individual’s WTP, in utils, for a given air purifier. We discuss this regression model further in the subsequent section. Because the two air purifiers differ only in their effectiveness at reducing indoor air pollution, we can use the difference in WTP across the two purifiers, scaled by the difference in their effectiveness

⁹A respondent who chooses an air purifier over \$10, then chooses \$20 over an air purifier, and then chooses an air purifier over \$30 exhibits “multiple-switching behavior”.

for a given respondent’s home, to recover each respondent’s marginal willingness to pay for a 1% reduction in air pollution at home.¹⁰

3.3 Second Follow-Up Survey

Six weeks after the first follow-up survey, we launched a second follow-up to understand whether any changes in beliefs induced by information provision persist outside of the 2-week window. This survey again asked respondents to answer Q1 through Q3. We did not measure willingness to pay in the long-run follow-up.

Experimenter Demand Effects

Experimenter demand effects are a standard concern in settings like ours, where an explicit information treatment may reveal the researcher’s expectations to respondents (De Quidt, Haushofer, and Roth, 2018). While we do not attempt to explicitly bound demand effects, our survey design reflects several of the strategies highlighted by Haaland, Roth, and Wohlfart (2023) for mitigating the impact of experimenter demand effects on outcomes of interest. Specifically, (a) we incentivize our initial belief measures and our follow-up measure of willingness-to-pay, (b) we frame the information provided in a neutral way, (c) we measure willingness-to-pay two weeks after providing information. Additionally, as pointed out by Haaland et al., our incentives may obfuscate the role of information provision from the respondents’ point of view. Because respondents know that their bonus depends on how close their response was to true values, respondents may not see the information provision as an explicit attempt to move their priors, but rather as information about how likely they are to receive a bonus.

3.4 Pre-Registration and Multiple-Hypothesis Testing

We pre-registered this experiment with the AEA RCT registry (*AEARCTR-0010435*). Prespecification presents a tradeoff for researchers: adding additional tests allows for additional credible discoveries, but adding tests can also reduce the power of the other hypotheses a researcher wishes to conduct due to adjustments for multiple-hypothesis testing. In our context, we have several outcomes – measures of misperception, belief updating, and willingness to pay for clean air – and several race and income subgroups where we could test for differences in these outcomes. To balance this tradeoff, we follow the procedure described by Anderson and Magruder (2017) and prespecify “main” hypotheses, which speak directly to our central research questions, and “auxiliary” hypotheses which answer other natural questions that come up in our setting. This procedure preserves statistical power for testing the main hypotheses and then uses a split-sample technique to efficiently test the auxiliary hypotheses.

Throughout the paper, we use a red dagger (†) to delineate tables that contain results that correspond to our pre-specified hypotheses. Each of these tables reports p -values corrected for multiple hypothesis testing using the Benjamini-Hochberg step-down adjustment (Benjamini and Hochberg, 1995). Appendix D contains additional details about these hypotheses, the split-sample approach, and the Benjamini-Hochberg procedure.

¹⁰Our survey elicits beliefs about outdoor air pollution but WTP for indoor air quality. Since PM_{2.5} readily penetrates buildings, and indoor concentrations are strongly correlated (Krebs, Burney, Zivin, and Neidell, 2021), we do not view this as a first-order concern.

4 Sample and Data

We conducted our survey through *Prolific*, an online survey platform designed for conducting academic and commercial research. In total, 5,231 individuals completed the initial survey, and 3,880 completed the follow-up survey. After applying several sample restrictions in the interest of eliminating low-quality responses, we are left with a panel of 3,804 participants.¹¹

Sample Demographics

Table 1 compares demographic characteristics of the respondents to the US population, and Appendix Figure A6 plots the spatial distribution of respondents across the United States. Our sample contains a somewhat lower fraction of Black and Latino respondents relative to the US population, and a somewhat higher fraction of Asian respondents relative to the US population, but otherwise matches the US population well.

The Appendix contains several additional tables that evaluate the balance and representativeness of our sample: Appendix Table A1 reports balance across treatment arms by race, and Appendix Table A2 reports balance by income group. Our treatment is balanced by income, and balanced for race groups with large numbers of respondents: White, Black or African, Asian, and 2 or more races. Appendix Table A3 calculates average particulate matter exposure by race using respondents’ reported ZIP codes. White respondents have the lowest pollution exposure ($7.75 \mu\text{g}/\text{m}^3$), and Asian respondents have the highest ($8.63 \mu\text{g}/\text{m}^3$). The Black-White pollution exposure gap is $0.49 \mu\text{g}/\text{m}^3$, which is similar to the population statistic reported by Currie, Voorheis, and Walker (2023). The average pollution level across all respondents is $7.95 \mu\text{g}/\text{m}^3$, which is also similar to estimates of the population average for the United States: The Air Quality Life Index, for example, estimated that population-weighted average particulate matter levels were $7.1 \mu\text{g}/\text{m}^3$ in 2020 and $8.2 \mu\text{g}/\text{m}^3$ in 2023 (Air Quality Life Index (AQLI), 2022, 2025).

Finally, Appendix Table A4 compares the sample demographics of the 5,137 respondents who were eligible to participate in the follow-up survey with the 3,804 participants who completed the follow-up survey to test whether attrition is correlated with observable characteristics. We find little evidence of attrition by race or pollution exposure, and our treatment remains balanced in the follow-up.¹²

¹¹We apply the following restrictions in order, omitting each response at the first condition it meets: (a) incomplete responses (401 from the initial survey and 3 from the follow-up survey); (b) duplicate IP addresses or duplicate Prolific accounts (11 duplicate IP addresses and 1 duplicate Prolific account); (c) responses completed in less than 90 seconds for the initial or follow-up survey (6 from the initial survey and 52 from the follow-up survey); and (d) responses that entered an invalid US ZIP code, or a ZIP code that did not match any ZIP code in our pollution dataset (76 responses). The respective median completion times for these surveys are 309 and 318 seconds.

¹²Note that the discrepancies in the race variables between the Table A4 (attrition) and Table 1 are the result of using Prolific’s automatically-reported respondent race variables rather than self-reported demographics. The latter were only collected in the follow-up survey, meaning we have to rely on Prolific’s measure when studying attrition by race. Similarly, we do not have information about the income of individuals who did not complete the follow-up survey.

Table 1: Sample Demographics

	Survey Sample	US Average	p-value
<i>Median HH Income</i>	\$65,000	\$69,021	0.365
<i>Median Age</i>	37	38.8	0.213
<i>Percent White</i>	75.8	75.8	0.957
<i>Percent Black</i>	8.5	13.6	<0.01
<i>Percent Asian</i>	9.2	6.1	<0.01
<i>Percent Latino</i>	12.2	18.9	<0.01

Notes: This table compares the demographics of the 3,804 respondents who completed both the initial and follow-up surveys to the demographics of the US population. The race and income summary statistics for the US population reflect data from the US Census 2022 Quick Facts; the median US age reflects data from the 2021 US Census Population and Housing Unit Estimates Tables. The final column contains a p -value from a one-sample test of difference in means between the survey sample and the US population.

Data on Beliefs and Choices

Our surveys provide us with data on respondents’ prior and posterior beliefs about local air pollution, as well as choices between each type of air purifier and gift cards of differing values.

We denote respondents’ beliefs $x_{i,p}^k$, where $p \in \{\text{pre}, \text{post}\}$ is the time period, and $k \in \{e, r, h\}$ indexes the topic: air pollution exposure in levels (e), relative air pollution exposure (r), or the health impact of their exposure (h). For each individual randomized into the treatment arm, we also have information shown to each individual ($s_{i,\text{shown}}^k$), as well as the information that they would have been shown ($s_{i,\text{alt}}^k$) in the alternative realization of the source-randomization.

The multiple price lists provide a total of 20 binary decisions for each individual, $c_{i,t}$, where i indexes the individual and t indexes the multiple price list row. In addition to individuals’ choices and the opportunity cost of each air purifier in each multiple price list row, we also have data on each air purifier’s effectiveness at reducing at-home air pollution for each individual: effectiveness_i^A and effectiveness_i^B . This is the only major attribute that differed between the two products.

5 Results

We begin by describing initial results surrounding misperceptions about air pollution and whether respondents update their beliefs in response to information. We then discuss the results of the multiple price list experiment and the derived measure of marginal willingness to pay for clean air.

5.1 Prior Beliefs

Figure 2 compares respondents’ beliefs about their relative pollution exposure — whether they believe they live in a clean or dirty part of their state — to the mean pollution estimates from Di et al. (2016) and Van Donkelaar et al. (2019). The 45-degree line plots the locus of perfectly-informed responses. If all respondents knew exactly where their neighborhood fell on the distribution of pollution in their state, responses would fall on this line. Instead, we see a trend in responses that is flatter: individuals in the least polluted ZIP codes in the state believe that their air pollution is in the 41st percentile of the state air quality distribution; individuals in the

most polluted neighborhoods in their state believe their air pollution is in the 62nd percentile.¹³

This pattern is consistent with significant misperceptions about relative air pollution exposure, but it does not suggest that beliefs about relative air pollution exposure are systematically optimistic or pessimistic. On average, respondents who live in ZIP codes with air pollution in the 50th percentile believed that their air pollution was slightly worse than the median ZIP code. Similar patterns have been observed elsewhere in the information provision literature: Cruces, Perez-Truglia, and Tetaz (2013), for example, find that people believe that they are closer to the middle of the income distribution than they are in reality. In a study of tax fairness, Nathan, Perez-Truglia, and Zentner (2024) similarly find that “individuals who pay more than average tend to underestimate how much they pay relative to others, and individuals who pay less than average tend to overestimate how much they pay.”


An important question is whether the patterns in Figure 2 indicate genuine misperception, or whether within-ZIP code variation in pollution leads us to conflate private information about exposure with misperception. The air pollution data provided by Di et al. (2016) and Van Donkelaar et al. (2019), however, suggest that there is little within-ZIP code variation in air pollution: 97.3% of variation in the remote-sensed air pollution data is explained by across-state and across-ZIP code variance. The remaining 2.7% of variation is a result of within-ZIP code variation in pollution levels.¹⁴

In Appendix Figure A11, we test whether the misperceptions in beliefs about relative air pollution exposure can be explained by long-term changes in air pollution exposure: Perhaps the observed misperceptions are a function of slow-to-update beliefs about pollution. We find little evidence to support this theory. Conditional on the air pollution percentile of a respondent’s ZIP code, there is no relationship between 10- or 20-year changes in ZIP code-level air pollution concentrations and respondent beliefs. In other words, consider two respondents living in different ZIP codes that are both in the 40th percentile of particulate air pollution today, but differed in air pollution levels in the past. The results in Figure A11 suggest that the 10- or 20-year changes in air pollution concentrations do not predict differences in misperceptions between these two individuals, on average.

¹³Note that the distribution of individuals across ZIP code pollution percentiles is not centered at 50 despite our survey pool having similar pollution exposure levels to the US population. This is likely a result of the unequal distribution of people across ZIP codes. Relative to rural ZIP codes, urban ZIP codes are both more polluted on average (Currie, Voorheis, and Walker, 2023), and more populous on average (Glimpse, 2014). As a result, if we measure pollution percentile at the ZIP code, a randomly-drawn individual from the US is more likely to live in a ZIP code with below median air quality than they are to live in a ZIP code with above median air quality.

¹⁴We also note that this is not a mechanical consequence of the resolution of these datasets. Both datasets are provided in 1 km^2 grids. The median ZIP code is 98.7 square kilometers (38.1 square miles), meaning that the typical ZIP code contains many pollution raster grid cells. Dense cities in the US tend to have smaller ZIP codes. Although very small ZIP codes may only feature a single-digit number of raster cells, average air pollution levels within small ZIP codes are unlikely to exhibit meaningful variation, given the dispersion patterns of particulate matter (See, for example, Garcia Rivera, Dinkelacker, Kioutsioukis, Adams, and Pandis (2021) regarding variation in pollution levels at differing raster sizes).

Figure 2: Initial Beliefs About Relative Air Pollution Exposure




Notes: This figure plots respondents' perceived relative air pollution exposure (their perception of their ZIP code's air pollution percentile) as a function of their measured relative air pollution exposure (the percentile of the mean of [Di et al. \(2016\)](#) and [Van Donkelaar et al. \(2019\)](#) for their ZIP code). Within each percentile bin, the black dot is the mean perception; the error bars contain 95% of responses. The dashed 45-degree line plots where perceptions would align with measurements from [Di et al. \(2016\)](#) and [Van Donkelaar et al. \(2019\)](#). The kernel density below the x-axis plots the distribution of respondents' within-state pollution percentile; the R-squared value corresponds to a simple bivariate regression of perceived pollution percentile on measured pollution percentile.

Figure 3 and Table 2 compare respondents' beliefs about their air pollution exposure in levels to estimates from these same two data sources. First, unlike beliefs about relative air pollution exposure, beliefs about air pollution exposure in levels appear biased upwards: on average, respondents believe that the air pollution levels where they live are higher than measurements suggest. Despite this bias in levels, the *slope* of beliefs roughly matches measurement — for every 1-unit increase in $PM_{2.5}$ exposure, beliefs about $PM_{2.5}$ exposure increase by 1.03 units.

Notably, the R-squared values that accompany Figures 2 and 3 show that beliefs about air pollution exposure (whether relative or absolute) explain very little of the variation in measured pollution. Beliefs about relative air pollution exposure, for example, explain 7% of measured variation in relative air pollution exposure; beliefs about air pollution levels explain roughly 1% of measured variation in air pollution levels.

Figure 3: Initial Beliefs About PM_{2.5} Pollution Exposure



Notes: This figure plots respondents' perceived PM_{2.5} air pollution exposure (their perception of their ZIP code's average PM_{2.5} pollution) as a function of academic estimates of their PM_{2.5} air pollution exposure (the mean of [Di et al. \(2016\)](#) and [Van Donkelaar et al. \(2019\)](#) for their ZIP code). Within each bin, the black dot is the mean perception; the error bars contain 95% of responses. The dashed line plots where perceptions would match measurements from Di et al. (2016) and van Donkelaar et al. (2019). The kernel density below the x-axis plots the distribution of respondents' PM_{2.5} pollution exposure; the R-squared value corresponds to a simple bivariate regression of perceived PM_{2.5} pollution exposure on measured PM_{2.5} pollution exposure.

Recall that in the survey, we provide a reference level to respondents — the Clean Air Act's 12 $\mu\text{g}/\text{m}^3$ National Ambient Air Quality Standard (NAAQS) — to help participants express their beliefs about their air pollution exposure levels in a unit of measurement they may not be entirely familiar with. This reference may have served as an anchor point for survey responses, leading to overall pessimistic beliefs (98% of individuals live in ZIP codes where PM_{2.5} lies below the Clean Air Act threshold). This being said, we believe several patterns in the data suggest anchoring alone cannot fully explain our findings.

First, if pure anchoring on the NAAQS were driving responses, we would expect beliefs to cluster tightly around 12 $\mu\text{g}/\text{m}^3$. Figure 3 shows substantial dispersion in beliefs, with many respondents reporting values well below or well above this threshold. Second, the slope of beliefs with respect to actual pollution is approximately one, suggesting that respondents do incorporate local information into their beliefs rather than simply anchoring on the reference point. If anchoring were the dominant force, we would expect a much flatter relationship between beliefs and actual pollution. Third, our elicitation of beliefs about *relative* pollution exposure—which was conducted without any numeric anchor—shows similar patterns of overestimation among respondents in relatively clean areas. Respondents in the cleanest decile of their state believe they are near the median (Figure 2). This convergent evidence across differently-anchored questions suggests that a general tendency toward pessimism about air quality is not merely an artifact of our survey design. Lastly, note that anchoring, if present, would bias our estimates of misperception but would not invalidate our experimental estimates of how information affects beliefs and WTP. The information provision generates variation in beliefs regardless of the source of initial misperceptions.

As an alternative, we can also use this reference point to provide a second and perhaps simpler way to explore the accuracy of individuals’ beliefs about pollution levels. Table 2 displays a 2-by-2 matrix where the rows are partitioned by whether an individual’s ZIP code is on average below the Clean Air Act NAAQS for $\text{PM}_{2.5}$ (“Attain”) or not (“Nonattain”). The columns are partitioned by individuals’ beliefs. While 98% of individuals live in ZIP codes where the average air pollution lies below the Clean Air Act threshold for particulate matter, 30% of individuals believe that the air pollution where they live exceeds this threshold. This comparison still suggests that beliefs are pessimistic, but it does not rely as heavily on numeric responses in a unit of measurement that is likely new to most respondents.


Table 2: Beliefs vs. Measurements, EPA $\text{PM}_{2.5}$ Nonattainment

	Believe Attain	Believe Nonattain
Attain	0.678	0.302
Nonattain	0.006	0.014

Notes: This table compares respondents’ beliefs about their particulate air pollution exposure to measurements from Di et al. (2016) and van Donkelaar et al. (2019). Respondents are categorized based on whether the average air pollution level in their ZIP code exceeds the Clean Air Act threshold of $12 \mu\text{g}/\text{m}^3$, and whether they believe their air pollution exposure exceeds this threshold.

Finally, Figure 4 compares respondents’ beliefs about how air pollution impacts life expectancy to the pollution-health concentration-response estimates from Anderson (2020) and Krewski et al. (2009). A large majority of respondents overestimate the pollution-health gradient: on average, respondents believed that a 50% increase in US air pollution levels would reduce average life expectancy in the US by roughly 2 years. This is an order of magnitude larger than the average of the gradients calculated by Anderson (2020) and Krewski et al. (2009).

Figure 4: Initial Beliefs About The Health Impacts of Air Pollution



Notes: This figure plots respondents’ perceptions about how average life expectancy would change if air pollution in the US were 50% higher than it is right now. Respondents were first asked whether they thought a 50% increase in air pollution would impact life expectancy in the US. If they responded “no,” they were coded as having a prior belief of zero for the pollution-life expectancy gradient. If they responded “yes,” they were shown a slider ranging from 0 to 4 years and asked to report their belief about how much a 50% increase in air pollution would decrease life expectancy. For reference, they were shown several other causes of mortality. The dotted line represents the average of the two academic studies (Anderson, 2020; Krewski et al., 2009) that we use during the information provision section of this survey.

Taken together, these results suggest that misperceptions about air pollution are widespread and substantial. We also find that misperceptions are idiosyncratic: Table A5 shows that individual-level errors across our three belief measures are largely uncorrelated. This suggests that misperceptions are not driven by a small subset of “misperceiving” types — individuals who hold inaccurate beliefs about their relative pollution exposure are not especially likely to also hold inaccurate beliefs about pollution levels or the health impacts of pollution.

We use respondent demographic data and estimates of misperception to test *Hypothesis Group 1: The accuracy of beliefs about air pollution is the same across race groups and income groups*. Table 3 compares misperceptions by income and race. We use an asterisk to denote results that are significant at the 0.05 level after accounting for multiple hypotheses, following Benjamini and Hochberg (1995). Appendix Table A21 contains the uncorrected p -values for each of our hypotheses. Across these six tests, we reject the null hypothesis only when comparing beliefs between Black and White respondents — specifically for beliefs about relative air pollution exposure and pollution exposure in levels. Here we find a modest difference in the accuracy of beliefs: The absolute error in beliefs about air pollution percentile is 3.7 percentile points smaller for White respondents than Black respondents. This difference persists even after conditioning on respondents’ measured air pollution percentile, suggesting it is not simply a product of Black and White respondents living in different parts of the pollution distribution. White respondents also held more accurate beliefs about pollution levels, meaning that they were less likely to overestimate their absolute air pollution exposure. We do not find significant differences between income groups for any of the three belief measures, nor do we find significant Black-White differences in beliefs about the health effects of air pollution.

Before presenting the results from the follow-up survey, it is worth briefly discussing how these findings relate to existing work. Often, economists posit that informational frictions will take the form of downward bias, where people underestimate pollution concentrations or the health impact of pollution. For example, [Hausman and Stolper \(2021\)](#) describe how environmental inequality can arise from a homogeneous scalar underestimation of pollution exposure; [Gao, Song, and Timmins \(2021\)](#) similarly model the difference between perceived and experienced pollution.¹⁵ Our findings are not consistent with this type of bias, and indeed suggest that to the extent that people in the US exhibit biased beliefs, they tend to overestimate both local air pollution concentrations and the air pollution-health gradient. We stress, however, that these findings should not be extrapolated beyond local air pollution. For example, in their paper, [Hausman and Stolper \(2021\)](#) discuss the case of toxic pollutants or novel chemicals with unknown health costs. In light of our findings here, we believe that studying misperceptions about other types of local pollutants may also provide useful contributions to the literature.

Table 3: Differences in Misperception by Race and Income (†)

Panel A: By Race

	(1) Pollution Percentile	(2) Pollution Level	(3) Life Expectancy
1(<i>White</i>)	−3.75* (0.97)	−4.33* (0.84)	−0.11 (0.05)
N	4066	4065	4066

Panel B: By Income

	(1) Pollution Percentile	(2) Pollution Level	(3) Life Expectancy
1(<i>Low Income</i>)	−0.67 (0.80)	0.29 (0.68)	0.07 (0.04)
N	2407	2407	2407

Notes: This table shows the differences in air pollution misperception between White and Black respondents (Panel A), and high- and low-income respondents (Panel B). High-income respondents are those with annual household income of over \$100,000 per year; low-income respondents are those with annual household income of less than \$50,000 per year. Respondents who are not either Black or White were not included in the Panel A regressions; middle-income respondents (those with annual incomes between \$50,000 and \$100,000 per year) were not included in the Panel B regressions. Each coefficient reflects a bivariate regression of individual-level absolute errors on a binary group indicator. In Panel A, a negative point estimate corresponds to lower levels of misperception for White respondents relative to Black respondents; in Panel B, a negative point estimate corresponds to lower levels of misperception for low-income respondents relative to high-income respondents. * Significant at the 0.05 level after accounting for multiple hypotheses, following [Benjamini and Hochberg \(1995\)](#).


¹⁵There is a large body of work in environmental economics, including [Gao, Song, and Timmins \(2021\)](#), that uses changes in China’s air pollution monitoring and disclosure to study the demand for clean air (see [Barwick, Li, Lin, and Zou, 2019](#); [Xie, Yuan, and Zhang, 2023](#); [Tu, Zhang, Xu, and Lu, 2020](#)). Existing models of bias in pollution beliefs may be informed by the fact that people in China tended to underestimate local air pollution exposure prior to the launch of a government monitoring program in 2013 ([Barwick et al., 2019](#)).

5.2 Posterior Beliefs

After the initial survey, three-fourths of respondents received a single piece of information corresponding to one of the three main survey questions. Two weeks later, we again asked respondents to share their beliefs. Here we describe respondents' beliefs in this follow-up survey, and whether we observe differences in updating between treated and control arms. Figures 5 through 7 compare the posterior beliefs between treated and untreated respondents. For simplicity, and to match subsequent regressions, each figure plots only 2 out of 4 total survey arms: the relevant “treated” arm and the pure control arm. For example, in the figure showing follow-up beliefs about relative air pollution exposure, we plot beliefs of individuals who received information about their ZIP code's air pollution percentile against those who received no information.

Figure 5 shows follow-up beliefs about relative pollution exposure for both the pure control group (red) and the treated group (blue). These results are consistent with partial belief updating: The beliefs of the treated group more closely track real-world measurements than do the control group, whose posteriors largely resemble the pattern of initial beliefs (see Figure 2).


Figure 5: Follow-Up Beliefs About Relative Air Pollution Exposure



Notes: This figure plots respondents' beliefs about their relative air pollution exposure during the follow-up survey. Perceptions for the control group are plotted in red; perceptions for the group that received information about the relative air quality in their ZIP code are plotted in blue. Perceptions about relative air quality for the two other treatment arms (air pollution levels and life expectancy) are not included in this figure. The dashed line is the locus where perceptions within an air pollution bin equal measured relative air pollution.

Figure 6 shows follow-up beliefs about pollution exposure in levels. In this case, there is not obvious evidence of belief updating. Both the pure control group and respondents who received information about the level of air pollution in their ZIP code show patterns similar to those in the initial survey. Specifically, beliefs about air pollution for the treated group remain pessimistic.¹⁶

Figure 6: Follow-Up Beliefs About Air Pollution Exposure




Notes: This figure plots respondents' beliefs about their air pollution exposure in levels during the follow-up survey. Perceptions for the control group are plotted in red; perceptions for the group that received information about the air pollution level in their ZIP code are plotted in blue. Perceptions about air quality for the two other treatment arms (air pollution levels and life expectancy) are not included in this figure. The dashed line is the locus where perceptions within an air pollution bin equal measured air pollution.

Figure 7 shows that respondents who received information about the pollution-mortality gradient strongly updated their priors. On average, the treated group (blue) believed that air pollution had a smaller impact on life expectancy than did the control group (red). This is consistent with respondents overestimating the pollution-health gradient in the initial survey, meaning that most of the individuals in the treated group received a signal that was lower than their prior.

¹⁶Note that in Figure 6, the majority of respondents have air pollution levels between 5 and 11 $\mu\text{g}/\text{m}^3$.

Figure 7: Follow-Up Beliefs About the Health Impacts of Air Pollution



Notes: This figure plots respondents' perceptions about the relationship between air pollution and average life expectancy reported during the follow-up survey. Perceptions for the control group are plotted in red; perceptions for the group that received information about the pollution-health gradient plotted in blue. Perceptions about the pollution-health gradient for the two other treatment arms are not included in this figure.

In Figures A12 through A14 of Appendix B, we investigate whether these changes in beliefs persist beyond our first 2-week follow-up survey. After 6 weeks, beliefs about the pollution-health gradient in the treated group were very similar to beliefs reported in the first follow-up. Beliefs about relative air pollution exposure at 6 weeks were also similar to beliefs at 2 weeks, but with some evidence of reversion toward the belief patterns exhibited by the control group. Overall, these patterns suggest that even relatively minor information interventions regarding air pollution can cause belief updating that is surprisingly persistent.

5.3 Updating Beliefs in Response to Information

Our aim in running this survey is to understand misperceptions about air pollution, how these misperceptions differ across demographic groups, and whether these misperceptions matter for willingness to pay for clean air. To answer the last question, we instrument for posterior beliefs using our randomized information provision. Here we present the first-stage results — i.e., the impact of information provision on posterior beliefs.¹⁷ These

¹⁷Recall that the information provided to individuals was source-randomized (i.e., pulled from one of two possible information sources). In Appendix A we also test whether respondents respond differently to the signal that they were shown versus the alternative signal they would have been shown in a different realization of the source-randomization. We conduct this test because source-randomization provides another potential instrument for follow-up beliefs. As discussed above, however, the simple information-provision instrument ($1[treated] \times (signal_i - prior_i)$) provides a stronger first stage than source-randomization alone; source randomization only provides a valid instrument for belief updating in one of our three arms. As such, we use the

regression results roughly correspond to Figures 5 through 7, above.

Updating: Signal versus Prior

We test whether respondents update their beliefs in response to each type of information provision using three regressions of the following form:

$$\underbrace{x_{i,post}^k - x_{i,prior}^k}_{\text{updating}} = \theta_0 + \theta_1^k \underbrace{z_i^k}_{\text{treat dummy}} \underbrace{(s_{i,signal}^k - x_{i,prior}^k)}_{\text{signal-prior}} + \theta_2^k (s_{i,signal}^k - x_{i,prior}^k) + \epsilon_i^k, \quad (6)$$

where k indexes the belief type: air pollution exposure in levels (e), relative air pollution exposure (r), and the health impact of pollution exposure (h). $x_{i,post}^k$ is individual i 's posterior belief about question k , and $x_{i,prior}^k$ is i 's prior about the same question. z_i^k is an indicator for whether or not individual i received information about question k .

θ_1^k , the coefficient on the interaction between the treatment and the belief gap, can be interpreted as Bayesian learning weights (Bottan and Perez-Truglia, 2022). These results are interesting per se, as they speak to whether it is easy to convince people to update incorrect priors, or if the misperceptions that we document persist even in the face of new information. A coefficient of one represents complete updating, whereas a coefficient of zero represents no updating. Finally, θ_2^k captures updating common to the treatment and control groups, which could reflect “spurious” learning, mean reversion or spotlight effects (Armantier, Nelson, Topa, Van der Klaauw, and Zafar, 2016; Cavallo, Cruces, and Perez-Truglia, 2017; Bottan and Perez-Truglia, 2025).

In all of the regressions that follow, we restrict the sample to include respondents that are part of the pure control group, or respondents where $z_i^k = 1$. For example, when estimating the impact of information provision on beliefs about relative pollution exposure, we include only individuals who received no information or those who received information about relative pollution exposure. We do so to avoid adjusting for cross-updating (i.e., an individual changing their beliefs about their relative air pollution exposure in response to information they received about another question).

Table 4 reports the results of estimating equation (6) for each of our three information treatments. Each column corresponds to a different arm of information provision; *Treated*(Signal — Prior)* is the coefficient of interest for each arm. Consistent with the figures presented above, these regressions suggest that treated respondents updated their beliefs in response to information about their *relative* air pollution exposure and the impact of air pollution on life expectancy.

The group that received information about their air pollution exposure in levels, however, did not update their beliefs relative to the control group. As discussed in Section 3, the fact that respondents internalize information about pollution ranks and health outcomes but not pollution levels may be explained by the fact that the latter requires understanding a new unit of measurement ($\mu\text{g}/\text{m}^3$). Percentile rankings and life-years, on the other hand, are potentially more familiar units that respondents could contextualize and remember (Chase and Simon, 1973; Bordalo, Gennaioli, and Shleifer, 2022).¹⁸

information-provision instrument in each of our three IV estimates of the impact of beliefs on MWTP for clean air.

¹⁸Figure A7 plots histograms of pollution level priors and posteriors. Respondents do not appear to bunch at the 12 $\mu\text{g}/\text{m}^3$

Finally, in each updating regression, we see evidence of updating common to the treatment and control groups ($\theta_2^k > 0$). This type of learning has been observed elsewhere in information provision settings (see Cullen and Perez-Truglia (2018) or Fuster, Perez-Truglia, Wiederholt, and Zafar (2022), for example). We are agnostic to the mechanism behind this learning, but we include a control for $(Signal - Prior)$ in our instrumental MWTP approach to disentangle responses to the treatment from common responses.

Table 4: Bayesian Belief Updating

	(1) Pollution Percentile	(2) Pollution Level	(3) Life Expectancy
$(Signal - Prior)$	0.29* (0.02)	0.77* (0.06)	0.44* (0.02)
$Treated*(Signal - Prior)$	0.31* (0.03)	0.04 (0.08)	0.30* (0.02)
N	1904	1914	1901

Notes: This table shows the results of applying equation (6) to each of the three questions we ask respondents: beliefs about relative air pollution exposure, beliefs about air pollution exposure in levels, and beliefs about the connection between air pollution and life expectancy. Each regression includes respondents from two of the four survey arms: the arm treated with the relative information, and the control group. *Signal* is the information that individuals either were shown or would have been shown had they been randomized into the treatment arm, *Prior* is an individual’s reported belief during the initial survey, and *Treated* is an indicator for whether an individual was randomized into the treatment arm. Standard errors are heteroskedasticity-robust; * denotes results that are significant at the 0.05 level. Note that the significance levels in this table are not corrected for multiple hypotheses, as the results presented in this table are descriptive and do not directly test any of our main hypotheses.

Differences in Updating

Many discussions of misperceptions about air pollution or other environmental harms take as given the idea that improving access to information about environmental harms will reduce inequality. This need not be the case depending on initial misperceptions, differences in updating, and differences in fully-informed willingness to pay.

With this motivation in mind, we tested for heterogeneity in updating (a) between Black and White respondents, and (b) between high- and low-income respondents as part of our auxiliary hypotheses. We separately tested for differences in updating for each of the three beliefs that we elicit: relative pollution exposure, pollution exposure in levels, and life expectancy. For each of these tests, we use the split-sample technique proposed in Anderson and Magruder (2017), using a 35% exploratory sample with a screening threshold of $t = 1.6$. As detailed in Section 3.4, this procedure is designed for settings where researchers have multiple hypotheses that differ in importance. The split-sample approach is used to test “auxiliary” hypotheses efficiently in order to preserve statistical power for the researchers’ main hypotheses.

We find relatively little evidence of differences in updating across income groups or between Black and White respondents. Across these 6 preliminary tests, only differences in updating between Black and White respondents with respect to life expectancy passes the screening threshold. Appendix D.2 contains the results of each of these exploratory regressions as well as the hold-out regressions for the single test that passed the screening threshold.


threshold, suggesting that the lack of updating may not reflect anchoring concerns related to the CAA threshold used as context in the survey question.

5.4 Willingness to Pay for Reductions in Air Pollution

The final set of survey results allows us to recover willingness to pay for reductions in air pollution. In the follow-up survey, each respondent filled out two multiple price lists; each price list asked the respondents to make 10 binary choices between gift cards of varying values and an air purifier. The two multiple price lists differed only in the attributes of the air purifier offered — one of the two air purifiers was more effective at reducing indoor air pollution. In this section, we first use the responses from these multiple price lists to recover individual-level estimates of WTP for each of the two air purifiers. The difference in willingness to pay for these two products together with the difference in effectiveness of the two purifiers allows us to recover each respondent’s willingness to pay for a 1 percent reduction in air pollution levels at home.

Figure 8 plots the demand curve for the two air purifiers. Each point on this graph represents the number of people who opted for a given air purifier over a gift card equal to price P . At every price point, demand for the higher-capacity air purifier (Air Purifier B) exceeds that of the low-capacity air purifier (Air Purifier A).

Figure 8: Demand for High- and Low-Capacity Air Purifiers



Notes: This figure plots the demand for air purifiers A and B, as implied by the multiple price list exercises conducted in the follow-up survey. For each point, the y-axis represents the gift-card price, and the x-axis represents the number of individuals who chose the air purifier over the gift card of that value. The marginal willingness to pay for clean air cannot be directly inferred from this graph because the difference in the “quality” of these air purifiers is a function of the size of the space that they clean, and therefore differs across individuals.

We estimate each individual’s WTP for each air purifier using a logit regression as proposed by [Jack, McDermott, and Sautmann \(2022\)](#). This approach models individual behavior by assuming that an agent chooses the option that maximizes her utility, subject to idiosyncratic choice errors which may lead to multiple-switching behavior. The latent utility of the air purifier for individual i in choice c of their multiple price list is:

$$V_{i,c} = \delta + \lambda p_{i,c} + u_i + \gamma_1 [\text{price order}_i] + \gamma_2 [\text{item order}_i] + \epsilon_{i,c} \quad (7)$$

where $p_{i,c}$ is the value of the gift card in choice c for individual i (i.e., the “price” of choosing the air purifier),

and δ is the value, in utils, of receiving an air purifier when the price equals zero. u_i is a fixed effect for each individual i . Assuming that $\epsilon_{i,c}$ is an independently and identically distributed extreme value type I error term yields the standard logit choice probabilities.

We include a vector of price order fixed effects [price order $_i$] and item order fixed effects [item order $_i$] to account for behavioral biases that may follow from the ordering of choices or columns in our multiple price list. These controls do not, however, address the problem of never-switchers. For individuals who never choose an air purifier within a given MPL, we code them as having a zero WTP for the corresponding air purifier.¹⁹ For individuals who always choose an air purifier within a given MPL, we code them as having a \$150 WTP. We discuss this assumption in [Appendix A.3](#) and show that our results are robust to alternative assumptions about the WTP of never-switchers. Appendix Table [A6](#) summarizes our willingness to pay estimates for each air purifier by demographic group.


From Equation ([Appendix A.3](#)) we have individual-level estimates of WTP for two air purifiers ($\frac{\delta+u_i^A}{\lambda}$ and $\frac{\delta+u_i^B}{\lambda}$, respectively). Let $effectiveness_i^A$ and $effectiveness_i^B$ be the effectiveness of each air purifier at reducing air pollution in the individual's home. This implies that individual i 's WTP for a 1% improvement in air quality is:

$$\widehat{MWP}_i = \frac{\widehat{WTP}_{B,i} - \widehat{WTP}_{A,i}}{effectiveness_{B,i} - effectiveness_{A,i}} \quad (8)$$

Figure [9](#) plots the CDF of willingness to pay. The mean WTP for a 1 percentage point decrease in air pollution at home is \$0.96. Using respondents' perceptions of their own air pollution exposure in levels, this corresponds to a willingness to pay of \$8 per $\mu\text{g}/\text{m}^3$.

¹⁹The smallest gift card has a \$1 value.

Figure 9: Marginal Willingness to Pay for Clean Air



Notes: This figure plots the CDF of estimates of MWTP from equation (8). We arrive at these estimates by converting the individual-level WTP for Air Purifier A and Air Purifier B into marginal willingness to pay for a one percentage point decrease in indoor air pollution. The individual-level WTP estimates for the two air purifiers reflect respondents' choices in the two multiple price list games during the follow-up survey. Because the information provided to respondents may impact MWTP differently for different groups, we only plot MWTP from the control group in this figure.

We also explore whether MWTP for clean air differs between Black and White respondents, or between high- and low-income respondents.

Table 5 displays the results from these two-sample t-tests. Additionally, Figures A9 and A10 plot the CDF of willingness to pay by race and income, respectively. Both tests in this hypothesis group reject the null that WTP is equal (including adjustments for multiple hypotheses). We find that White respondents are willing to pay 26 cents more to reduce pollution by 1% when compared to Black respondents, and high-income respondents are willing to pay 44 more cents to reduce pollution by 1% than are low-income respondents. It is relatively straightforward to translate these into percent differences, as the unconditional mean MWTP is about \$1 for a 1% reduction in pollution.

Given the gap in income between Black and White households in the US, it is natural to ask whether the difference in MWTP shown in column 1 of Table 5 is a direct consequence of differences in income across groups. In Appendix Table A8, we re-estimate the difference in MWTP between Black and White respondents while controlling for income. Specifically, we include fixed effects for each of the 16 self-reported income bins. The resulting coefficient reflects differences in MWTP between Black and White respondents within relatively narrow income bands. The conditional gap in MWTP is about 25% smaller than the unconditional gap, but this gap remains substantial — it is roughly 20% of the mean MWTP for clean air.²⁰

²⁰Note that estimating the gap in MWTP conditional on income was not pre-specified; this is an ad hoc test that we performed after estimating the initial set of results.

Table 5: Differences in Willingness to Pay for Clean Air by Group †

	White vs. Black	Low vs. High Income
1(<i>White</i>)	0.27* (0.08)	
1(<i>Low Income</i>)		-0.44* (0.06)
Num. obs.	3203	2405

Notes: This table displays the results of two bivariate regressions used to compare MWTP for clean air between Black and White respondents (column 1) and to compare between high- and low-income individuals (columns 2). We define “high-income” individuals as those who make more than \$100,000 annually. Low-income individuals are those who make less than \$50,000 annually. *Significant at the 0.05 level after accounting for multiple hypotheses.

Lastly, we ask whether individuals’ beliefs about air pollution (relative exposure, exposure in levels, and the relationship between pollution and health) impact their marginal willingness to pay for clean air. We test these hypotheses using an instrumental variables approach. We run three separate regressions, each comparing marginal willingness to pay between the control group and a single treatment group:

$$MWTP_i = \alpha + \beta^k(\widehat{posterior}_i^k) + \gamma(signal_i^k - prior_i^k) + \epsilon_i \quad (9)$$

where k indexes the belief type: life expectancy, pollution level, pollution percentile, and $posterior_i^k$ is i ’s follow-up belief about issue k . We use our information provision experiment to instrument for $posterior_i^k$. Specifically, the excluded instrument is the interaction between a treatment indicator and the difference between signal and prior.²¹ We include $signal_i^k - prior_i^k$ for all individuals to control for spurious updating.

Table 6 presents the estimates. We do not find evidence that updating beliefs about one’s *relative* pollution exposure impacts marginal willingness to pay for clean air and can rule out large elasticities.²² We find, however, that changing beliefs about the relationship between life expectancy and air pollution exposure increases respondents’ willingness to pay for reductions in air pollution. As beliefs about the impact of pollution on life expectancy increase by 1 year, willingness to pay for a 1 percent reduction in air pollution increases by 28 cents (or roughly 28%). Because the majority of respondents overestimated the pollution-health gradient, this coefficient largely reflects participants who both revise downwards their priors about the health impacts of pollution and lower their MWTP for clean air. Figure A8 shows that the inverse also holds: among respondents who initially *underestimated* the health costs of air pollution, those who received information about health costs *increased* their MWTP for clean air relative to the control group.

²¹The first-stage regression equation is $posterior_i^k = \pi_0 + \pi_1 z_i^k (signal_i^k - prior_i^k) + \pi_2 (signal_i^k - prior_i^k) + \eta_i$ where z_i^k is treatment dummy that equals one if and only if individual i received information about k . For individuals who did not receive treatment, $signal_i^k$ is defined as the information that individual i would have been shown had they been in the treated group.

²²Recall that respondents did not respond strongly to information about their air pollution exposure in *levels*, and thus the corresponding first stage F-statistic is weak.

Table 6: Estimating the Relationship Between Pollution Beliefs and Marginal Willingness to Pay for Clean Air [†]

	(1) MWTP	(2) MWTP	(3) MWTP
<i>Percentile beliefs</i>	0.003 (0.005)		
<i>Pollution level beliefs</i>		0.194 (0.436)	
<i>Life expectancy beliefs</i>			0.282* (0.104)
First-Stage F	196.719	0.630	226.141
N	1904	1912	1901

Notes: This table displays the results of equation (9), an instrumental variables regression which estimates the relationship between beliefs and marginal willingness to pay (MWTP) for a 1% reduction in air pollution at home. In all regressions, the signal shown to respondents acts as an instrument for posterior beliefs. Each regression includes respondents from two of the four survey groups: the pure control group and the group treated with a given type of information (i.e., relative pollution exposure, pollution exposure in levels, or life expectancy). For reference, the mean MWTP in our sample is roughly \$1 for a 1% reduction in air pollution at home. Heteroskedasticity-robust standard errors are in parentheses. * Significant at the 0.05 level after accounting for multiple hypotheses.

How do these results compare to other estimates of MWTP for PM_{2.5} improvements from the literature? There are surprisingly few existing studies to act as points of comparison. Most modern estimates of MWTP for air quality improvements are from papers set outside of the United States. Additionally, the few related papers recover WTP for related but not identical local pollutants, such as PM₁₀ or total suspended particulates. As such, we need to convert these estimates prior to making direct comparisons. For example, Ito and Zhang (2020) find that households in China are willing to pay \$1.34 per year to remove 1 $\mu\text{g}/\text{m}^3$ of PM₁₀ from the air. While there is no perfect way to convert PM₁₀ to PM_{2.5}, a common rule-of-thumb ratio used internationally, including by the World Health Organization, is: $\text{PM}_{2.5} \approx 0.6 \times \text{PM}_{10}$. Using this conversion suggests the MWTP from Ito and Zhang is equivalent to \$0.80 per unit of PM_{2.5}. In inflation-adjusted, purchasing power parity terms,²³ this point estimate is \$1.72 per year (2023\$). The authors suggest the average lifetime of an air filter to be 5 years, and thus MWTP for a 5-year reduction in PM_{2.5} would be around \$8.6, an estimate quite close to the mean MWTP we recover.

Relatedly, Freeman, Liang, Song, and Timmins (2019) use a residential sorting model to recover MWTP for PM_{2.5} improvements. They find that the typical household in China would be willing to pay \$21.70 (2005\$) for a 1 $\mu\text{g}/\text{m}^3$ reduction in annual average PM_{2.5} concentrations. Converting this to 2023 dollars using estimates of Chinese annual inflation yields a MWTP of \$30.38, which is equivalent to \$54.25 in terms of purchasing power parity. This estimate reflects MWTP for a *permanent* reduction in air pollution. If we were to convert our estimate to a permanent reduction using a 5 percent annual discount rate, our results suggest that respondents would be willing to pay \$29.21 for a permanent 1 $\mu\text{g}/\text{m}^3$ reduction in PM_{2.5}.²⁴ Even so, it is difficult to understand in both studies whether the study participants are fully informed and the effect that this may have on the authors' WTP estimates.

²³Specifically, we convert the 2014 estimate into 2023 dollars using a 2 percent annual inflation rate for China over this period, and then account for the differences in purchasing power to translate the 2023 value into MWTP for U.S. consumers.

²⁴With a 5% discount rate, a willingness to pay of \$8 spread evenly over five years (\$1.60 per year) translates into a lifetime willingness to pay (over 50 years) of approximately \$29.21.

5.5 Counterfactual MWTP under full belief updating

What do these results imply for the difference between WTP under the status quo and WTP with fully-informed consumers? Recall that in the follow-up survey, belief updating was incomplete, meaning that WTP in the treated groups does not reflect the WTP of individuals who are fully informed about a given aspect of pollution exposure. To estimate a counterfactual where our respondents update fully, we can start with the estimates of WTP from our control group as a baseline and then use our results above to extrapolate to the WTP of a fully informed individual.

For example, Table A12 shows that the control group has an average WTP of \$1.09 for a 1% improvement in air quality at home. In our control group, the average respondent overestimates the life expectancy impacts of a 50% increase in air pollution by 1.91 years. Column 3 of Table 6 gives the gradient between beliefs about the health effects of pollution and marginal willingness to pay for air quality improvements; for every additional year increase in the perceived pollution-life expectancy gradient, MWTP for clean air at home increases by \$0.28. Applying this gradient to the belief gaps in our control group implies that if the respondents in our survey were fully informed about the relationship between life expectancy and air pollution, their MWTP would be \$0.56 for a 1% improvement in air quality at home – 49% lower than the status quo.

Although we do not find a statistically significant relationship between beliefs about relative air pollution exposure and demand for clean air, we can use our point estimate and confidence interval around this estimate to perform a similar exercise.²⁵ It is worth noting that the point estimate of the effect of pollution percentile beliefs on willingness to pay (column 1 of Table 6) is “perverse”, albeit small. The interpretation of this coefficient is that when a respondent updates their belief about their local air quality in the positive direction by 1 percentile point (e.g., changing their belief about their pollution percentile from the 50th percentile of air quality to the 51st percentile of air quality), their willingness to pay for clean air at home increases by \$0.003 (i.e., they are willing to pay an additional 0.3 cents for a 1% increase in air quality at home). Taking this point estimate at face value, we can again use it as a gradient to infer how MWTP would change if the control group were to fully update about their relative pollution exposure. In this case, the implied change in WTP is small — rectifying misperceptions would reduce MWTP for a 1% improvement in home air quality from \$1.09 to \$1.04.

The standard errors around the slope between MWTP and pollution percentile beliefs are also useful for counterfactual exercises. While beliefs about relative air pollution are not severely optimistic or pessimistic, on average, respondents think their relative air quality is better than measurements suggest (see column 1 of Table A12). This fact suggests that if people respond to information about their relative air pollution ranking as one might *ex ante* expect – i.e., willingness to pay increases when you learn that your air quality ranking is worse than expected – then information about relative air pollution exposure might act in the opposite direction of information about the health impacts of air pollution. The confidence intervals around our point estimate allow us to speak to how large such an increase in WTP might be. Specifically, we can reject large increases in marginal willingness to pay in a world where respondents hold correct beliefs about their relative air pollution exposure. For example, using the lower bound of our 95% confidence interval for the estimate of the slope between beliefs about relative air pollution and the demand for clean air, correcting beliefs about relative air

²⁵We do not perform this exercise for pollution *levels* because the lack of a first stage means that our point estimate and standard errors are unlikely to be informative.

pollution exposure would increase mean marginal willingness to pay by just 11%, on average.²⁶

In summary, information in this environment serves primarily to correct overly pessimistic beliefs. On average, fully informed respondents lower their willingness to pay for air quality as compared to the uninformed or partially informed status quo.

5.6 Implications for environmental inequality

What do these findings mean for measures of environmental inequality or gaps in air pollution exposure? We find differences in Black and White respondents' willingness to pay for clean air, even conditional on income. Our descriptive results provide some suggestive evidence that misperception could play a role in the unexplained portion of this willingness to pay gap: As noted above, Black respondents hold less accurate beliefs about their pollution exposure in levels and their relative pollution exposure than White respondents.

At the same time, however, we do not find strong evidence that rectifying these types of misperceptions leads to changes in willingness to pay for clean air. As such, the willingness to pay gap persists even under optimistic counterfactuals. Specifically, if we (a) ignore misperceptions about the pollution-health gradient, and (b) take the 95th percentile gradient between relative pollution beliefs and WTP, then correcting misperceptions would only shrink the Black-White WTP gap by roughly one quarter to one third.²⁷

Correcting misperceptions about the relationship between air pollution and health is unlikely to close the Black versus White gap in willingness to pay for clean air; our results suggest that doing so may increase differences in WTP. If we assume a common slope in beliefs (i.e., the coefficient from column 3 of Table 6), marginal willingness to pay for clean air would fall by roughly the same amount for both Black and White individuals in levels, leading to an increase in the gap in proportional terms. If we allow for differences in the relationship between MWTP and beliefs by race, we find that the gradient between life expectancy and marginal willingness to pay for clean air between Black and White respondents would fall in raw terms, but still increase in proportional terms.

Overall, our survey suggests that even though some differences exist in the accuracy of beliefs across groups, these differences are modest, and it seems unlikely that information provision alone would improve the Black versus White gaps in willingness to pay for clean air or experienced exposure levels.

6 Conclusion

Environmental inequality is one of the most durable features of the American environmental landscape. Understanding why it persists requires distinguishing among competing hypotheses—differences in income, heterogeneous preferences, discriminatory institutions, and information frictions—that have proven difficult to disentangle with existing data. A key obstacle is that revealed preference data alone do not allow researchers to distinguish between beliefs and preferences, making it impossible to determine whether demographic gaps in

²⁶This “lower bound” slope is -0.007. In words, this coefficient means that when a respondent updates their belief about their local air quality in the positive direction by 1 percentile point, their willingness to pay for a 1% increase in air quality at home decreases by 0.7 cents.

²⁷Using the lower bound of the 95% confidence interval (-0.007) as the slope between relative pollution beliefs and WTP, the unconditional White-Black gap falls from \$0.268 to \$0.203 (24.3%) and the conditional gap falls from \$0.182 to \$0.117 (35.7%).

WTP reflect differences in what people want or differences in what they know. This paper uses an incentivized survey experiment to address this challenge, providing what we believe is the first empirical decomposition of observed WTP gaps for air quality into components attributable to differential information and residual preference heterogeneity.

Our findings speak to several open questions in the environmental justice and environmental valuation literatures. We organize our conclusions around the four hypotheses we set out to test.

First, do Americans hold accurate beliefs about air pollution? They do not. Respondents overestimate both local pollution levels and the health consequences of exposure by roughly an order of magnitude. Beliefs about relative exposure are unbiased on average but exhibit substantial compression—residents of the cleanest areas think they face more pollution than they do, while residents of the dirtiest areas think they face less. These misperceptions are large enough to meaningfully affect revealed-preference measures of willingness to pay: our counterfactual exercises suggest that fully informed MWTP would be roughly half of what we observe under existing beliefs.

Second, do misperceptions differ across the demographic groups that are central to environmental justice debates? Largely, no. We find similar beliefs about pollution levels and health impacts across racial and income groups. Across six tests for differences in belief accuracy by race and income, we reject the null in only two cases, both comparing White and Black respondents: White respondents hold modestly more accurate beliefs about their relative exposure and their exposure in levels. These differences are limited in their implications for WTP, because beliefs about pollution exposure do not significantly affect MWTP in our estimates. This similarity of beliefs across groups implies that information frictions, at least as they operate through the channels we measure, cannot be a primary driver of environmental inequality in the United States.

Third, does correcting misperceptions causally affect demand for clean air? Yes, but only through a specific channel. Beliefs about the health impacts of pollution are the primary driver of MWTP: exogenously shifting perceived health impacts has a large, precisely estimated effect on WTP. Beliefs about relative exposure do not significantly affect MWTP, even though our information treatment shifts them. Similarly, we do not find evidence that MWTP is affected by beliefs about pollution levels, with the caveat that information does not strongly impact beliefs about pollution levels, meaning that the first-stage in this regression is weak.

Fourth, would correcting misperceptions close observed gaps in willingness to pay between demographic groups? Our counterfactual exercises suggest not. Because misperceptions are similar across groups, correcting them shifts everyone’s WTP in the same direction, leaving relative gaps largely unchanged. This constitutes direct evidence against the hypothesis that information frictions are a primary sustaining force behind environmental inequality, at least through the belief-to-WTP channel.

The persistence of WTP gaps after experimentally removing information differences has a constructive implication: it allows us to characterize the nature of these gaps more precisely. The residual WTP differences that survive belief correction—higher MWTP among White and high-income respondents, even when all groups hold accurate beliefs—more closely approximate differences in underlying preferences than do raw observed WTP gaps. To our knowledge, these “informed” estimates of WTP heterogeneity across race and income are the first of their kind for air quality in the United States, and provide a new empirical input for models of environmental sorting and policy design that require preference parameters by demographic group. The existence of within-

income gaps in WTP that persist after correcting beliefs suggests that income constraints alone cannot fully explain the preference gap, and that other factors—the amenity bundles with which environmental quality is packaged, differences in baseline health, or taste heterogeneity—may also play a role.

Several caveats are worth noting. Our analysis focuses on the belief-to-WTP channel, but environmental inequality also operates through political engagement, collective action, and the supply side of pollution. Information frictions could matter through these channels even if they do not significantly affect individual WTP. A complete accounting of how misperceptions shape environmental inequality would require understanding how beliefs interact with civic participation and the political economy of environmental regulation, which we view as a fruitful direction for future research. Finally, our findings are specific to the US context, where air quality is relatively good, and beliefs tend to be overly pessimistic. In settings with higher pollution levels where people underestimate exposure, the information channel may play a more prominent role in sustaining disparities, and correcting beliefs would be expected to *increase* rather than decrease WTP.

References

- Air Quality Life Index (AQLI). 2022. “United States Fact Sheet.” Tech. rep., Energy Policy Institute at the University of Chicago (EPIC). URL <https://aqli.epic.uchicago.edu>.
- . 2025. “United States Fact Sheet.” Tech. rep., Energy Policy Institute at the University of Chicago (EPIC). URL https://aqli.epic.uchicago.edu/files/USFactSheet_2025.pdf.
- Anderson, Michael L. 2020. “As the wind blows: The effects of long-term exposure to air pollution on mortality.” *Journal of the European Economic Association* 18 (4):1886–1927.
- Anderson, Michael L and Jeremy Magruder. 2017. “Split-sample strategies for avoiding false discoveries.” Tech. rep., National Bureau of Economic Research.
- Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert Van der Klaauw, and Basit Zafar. 2016. “The price is right: Updating inflation expectations in a randomized price information experiment.” *Review of Economics and Statistics* 98 (3):503–523.
- Banzhaf, Spencer, Lala Ma, and Christopher Timmins. 2019. “Environmental justice: The economics of race, place, and pollution.” *Journal of Economic Perspectives* 33 (1):185–208.
- Barwick, Panle Jia, Shanjun Li, Liguang Lin, and Eric Zou. 2019. “From fog to smog: The value of pollution information.” Tech. rep., National Bureau of Economic Research.
- Been, Vicki and Francis Gupta. 1997. “Coming to the nuisance or going to the barrios? A longitudinal analysis of environmental justice claims.” *Ecology law quarterly* 24 (1):1–56.
- Benjamini, Yoav and Yosef Hochberg. 1995. “Controlling the false discovery rate: a practical and powerful approach to multiple testing.” *Journal of the Royal statistical society: series B (Methodological)* 57 (1):289–300.


- Benjamini, Yoav and Daniel Yekutieli. 2001. "The control of the false discovery rate in multiple testing under dependency." *Annals of statistics* :1165–1188.
- Bhuller, Manudeep and Henrik Sigstad. 2022. "2SLS with multiple treatments." *arXiv preprint arXiv:2205.07836* .
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2022. "Salience." *Annual Review of Economics* 14 (1):521–544.
- Bottan, Nicolas and Ricardo Perez-Truglia. 2025. "Betting on the house: Subjective expectations and market choices." *American Economic Journal: Applied Economics* 17 (1):459–500.
- Bottan, Nicolas L and Ricardo Perez-Truglia. 2022. "Choosing your pond: location choices and relative income." *Review of Economics and Statistics* 104 (5):1010–1027.
- Bullard, Robert D. 1993. "Race and environmental justice in the United States." *Yale J. Int'l L.* 18:319.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation expectations, learning, and supermarket prices: Evidence from survey experiments." *American Economic Journal: Macroeconomics* 9 (3):1–35.
- Chase, William G and Herbert A Simon. 1973. "Perception in chess." *Cognitive psychology* 4 (1):55–81.
- Christensen, Peter and Christopher Timmins. 2022. "Sorting or steering: The effects of housing discrimination on neighborhood choice." *Journal of Political Economy* 130 (8):2110–2163.
- Colmer, Jonathan M, Suvy Qin, John L Voorheis, and Reed Walker. 2026. "Income, wealth, and environmental inequality in the united states." Tech. rep., National Bureau of Economic Research.
- Cruces, Guillermo, Ricardo Perez-Truglia, and Martin Tetaz. 2013. "Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment." *Journal of public Economics* 98:100–112.
- Cullen, Zoë and Ricardo Perez-Truglia. 2018. "How much does your boss make? the effects of salary comparisons." Tech. rep., National Bureau of Economic Research.
- Currie, Janet, John Voorheis, and Reed Walker. 2023. "What caused racial disparities in particulate exposure to fall? New evidence from the Clean Air Act and satellite-based measures of air quality." *American Economic Review* 113 (1):71–97.
- De Quidt, Jonathan, Johannes Haushofer, and Christopher Roth. 2018. "Measuring and bounding experimenter demand." *American Economic Review* 108 (11):3266–3302.
- Depro, Brooks, Christopher Timmins, and Maggie O'neil. 2015. "White flight and coming to the nuisance: can residential mobility explain environmental injustice?" *Journal of the Association of Environmental and resource Economists* 2 (3):439–468.

- Di, Qian, Itai Kloog, Petros Koutrakis, Alexei Lyapustin, Yujie Wang, and Joel Schwartz. 2016. “Assessing PM_{2.5} exposures with high spatiotemporal resolution across the continental United States.” *Environmental science & technology* 50 (9):4712–4721.
- Fowlie, Meredith, Edward Rubin, and Reed Walker. 2019. “Bringing satellite-based air quality estimates down to earth.” In *AEA Papers and Proceedings*, vol. 109. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 283–288.
- Freeman, Richard, Wenquan Liang, Ran Song, and Christopher Timmins. 2019. “Willingness to pay for clean air in China.” *Journal of Environmental Economics and Management* 94:188–216.
- Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar. 2022. “Expectations with endogenous information acquisition: An experimental investigation.” *Review of Economics and Statistics* 104 (5):1059–1078.
- Gao, Xuwen, Ran Song, and Christopher Timmins. 2021. “The role of information in the rosen-roback framework.” Tech. rep., National Bureau of Economic Research.
- Garcia Rivera, Pablo, Brian T Dinkelacker, Ioannis Kioutsoukis, Peter J Adams, and Spyros N Pandis. 2021. “Source-resolved variability of fine particulate matter and human exposure in an urban area.” *Atmospheric Chemistry and Physics Discussions* 2021:1–36.
- Glimpse, Warren. 2014. “ZIP Code Urban/Rural Demographic Patterns.” URL https://proximityone.wordpress.com/2014/10/11/zip-code-urbanrural-demographic-patterns/?utm_source=chatgpt.com. Accessed: 2025-10-24.
- Greenstone, Michael and Claire Qing Fan. 2018. “Introducing the air quality life index.” *AQLI Annual Report* .
- Grossman, Michael. 1972. “On the concept of health capital and the demand for health.” *Journal of Political economy* 80 (2):223–255.
- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart. 2023. “Designing information provision experiments.” *Journal of economic literature* 61 (1):3–40.
- Hausman, Catherine and Samuel Stolper. 2021. “Inequality, information failures, and air pollution.” *Journal of Environmental Economics and Management* 110:102552.
- Ito, Koichiro and Shuang Zhang. 2020. “Willingness to pay for clean air: Evidence from air purifier markets in China.” *Journal of Political Economy* 128 (5):1627–1672.
- Jack, B Kelsey, Kathryn McDermott, and Anja Sautmann. 2022. “Multiple price lists for willingness to pay elicitation.” *Journal of Development Economics* 159:102977.
- Krebs, Benjamin, Jennifer Burney, Joshua Graff Zivin, and Matthew Neidell. 2021. “Using crowd-sourced data to assess the temporal and spatial relationship between indoor and outdoor particulate matter.” *Environmental science & technology* 55 (9):6107–6115.

- Krewski, Daniel, Michael Jerrett, Richard T Burnett, Renjun Ma, Edward Hughes, Yuanli Shi, Michelle C Turner, C Arden Pope III, George Thurston, Eugenia E Calle et al. 2009. *Extended follow-up and spatial analysis of the American Cancer Society study linking particulate air pollution and mortality*. 140. Health Effects Institute Boston, MA.
- Kuminoff, Nicolai V, V Kerry Smith, and Christopher Timmins. 2013. “The new economics of equilibrium sorting and policy evaluation using housing markets.” *Journal of economic literature* 51 (4):1007–1062.
- Meng, Jun, Chi Li, Randall V Martin, Aaron Van Donkelaar, Perry Hystad, and Michael Brauer. 2019. “Estimated long-term (1981–2016) concentrations of ambient fine particulate matter across North America from chemical transport modeling, satellite remote sensing, and ground-based measurements.” *Environmental science & technology* 53 (9):5071–5079.
- Nathan, Brad C, Ricardo Perez-Truglia, and Alejandro Zentner. 2024. “Paying your fair share: Perceived fairness and tax compliance.” Tech. rep., National Bureau of Economic Research.
- Pastor, Manuel, Jim Sadd, and John Hipp. 2001. “Which came first? Toxic facilities, minority move-in, and environmental justice.” *Journal of urban affairs* 23 (1):1–21.
- Tu, Meng, Bing Zhang, Jianhua Xu, and Fangwen Lu. 2020. “Mass media, information and demand for environmental quality: Evidence from the “Under the Dome”.” *Journal of Development Economics* 143:102402.
- Van Donkelaar, Aaron, Randall V Martin, Chi Li, and Richard T Burnett. 2019. “Regional estimates of chemical composition of fine particulate matter using a combined geoscience-statistical method with information from satellites, models, and monitors.” *Environmental science & technology* 53 (5):2595–2611.
- Wolverton, Ann. 2009. “Effects of socio-economic and input-related factors on polluting plants’ location decisions.” *The BE Journal of Economic Analysis & Policy* 9 (1).
- Xie, Tingting, Ye Yuan, and Hui Zhang. 2023. “Information, awareness, and mental health: Evidence from air pollution disclosure in China.” *Journal of Environmental Economics and Management* 120:102827.

Appendix Figures


Figure A1: Initial Survey




Notes: 5,231 individuals completed the initial survey on the platform *Prolific*. Respondents were eligible if they were over 18 years old and lived in the continental US.

Figure A2: Satellite-Based $\text{PM}_{2.5}$ Measures

(a) Di et al. (2016)




(b) van Donkelaar et al. (2019)




Notes: This figure plots three-year average $\text{PM}_{2.5}$ estimates based on [Di et al. \(2016\)](#) and [Van Donkelaar et al. \(2019\)](#). Both of these studies use satellite-derived measures of aerosol optical depth (AOD) together with air pollution monitoring data to build models that convert AOD data and other observable inputs into estimates of particulate matter ($\text{PM}_{2.5}$) concentrations. The model-derived outputs consist of a 1 km by 1 km raster of the estimated average annual $\text{PM}_{2.5}$ concentrations.

Figure A3: Source-Randomized PM_{2.5} Pollution Information



Notes: This figure plots source randomized information about relative PM_{2.5} pollution (right panel) and PM_{2.5} levels (left panel). Each point in this figure represents a ZIP code; the x-axis is the information signal derived from [Van Donkelaar et al. \(2019\)](#), and the y-axis is information signal derived from [Di et al. \(2016\)](#).

Figure A4: Follow-up survey



Notes: In the follow-up survey, respondents were asked the same three questions about air pollution before completing two incentivized multiple price list exercises designed to elicit willingness to pay for improvements in air quality.


Figure A5: Example of a Multiple Price List used for Eliciting WTP

*There are 10 pairs of items below. For each case, choose the item that you would prefer to receive. Recall that **Air Purifier B** would reduce indoor air pollution in your home by **51%***

Choice 1	Air Purifier B	\$1 Amazon Gift Card
Choice 2	Air Purifier B	\$5 Amazon Gift Card
Choice 3	Air Purifier B	\$15 Amazon Gift Card
Choice 4	Air Purifier B	\$25 Amazon Gift Card
Choice 5	Air Purifier B	\$35 Amazon Gift Card
Choice 6	Air Purifier B	\$45 Amazon Gift Card
Choice 7	Air Purifier B	\$60 Amazon Gift Card
Choice 8	Air Purifier B	\$75 Amazon Gift Card
Choice 9	Air Purifier B	\$95 Amazon Gift Card
Choice 10	Air Purifier B	\$130 Amazon Gift Card


Notes: Each respondent in the follow-up survey was shown two multiple price lists: one for the high-capacity air purifier, and one for the low-capacity air purifier. The description of the air purifier's effectiveness is based on the respondent's home size and the purifier's clean air delivery rate. The order of columns (air purifier first or gift cards first) was randomized across individuals, as was the direction of the gift card column (ascending or descending).

Figure A6: Respondent Locations




Notes: This figure plots the location (ZIP code centroid) of the 3,804 respondents who completed both the initial and the follow-up surveys.

Figure A7: Histograms of Prior and Posterior Beliefs about Pollution Levels




Notes: This figure plots histograms for prior (left panel) and posterior (right panel) beliefs about $\text{PM}_{2.5}$ levels. The vertical dotted line in each figure is at $12\mu\text{g}/\text{m}^3$, which was the reference level provided to respondents. In both figures, the 50+ bin contains counts for all responses of $50\mu\text{g}/\text{m}^3$ or more.

Figure A8: MWTP by Initial Misperception (Life Expectancy)




Notes: This figure plots MWTP for clean air against participants' initial misperceptions about the impacts of a 50% increase in local air pollution on average life expectancy. The treated group received information about the impacts of air pollution on life expectancy prior to MWTP elicitation; the control group did not receive any such information. Participants from the other two survey arms (information about relative pollution exposure and information about pollution exposure in levels) are not included in this figure. The MWTP measure reflects a participant's demand for a 1% reduction in air pollution levels at home.

Figure A9: Marginal Willingness to Pay for Clean Air by Race




Notes: This figure plots the CDF of estimates of MWTP from equation (8), by race. We arrive at these estimates by converting the individual-level WTP for Air Purifier A and Air Purifier B into marginal willingness to pay for a one percentage point decrease in indoor air pollution. The individual-level WTP estimates for the two air purifiers reflect respondents' choices in the two multiple price list games during the follow-up survey. Because the information provided to respondents may impact MWTP differently for different groups, we only plot MWTP from the control group in this figure.

Figure A10: Marginal Willingness to Pay for Clean Air by Income



Notes: This figure plots the CDF of estimates of MWTP from equation (8) by income group. We arrive at these estimates by converting the individual-level WTP for Air Purifier A and Air Purifier B into marginal willingness to pay for a one percentage point decrease in indoor air pollution. The individual-level WTP estimates for the two air purifiers reflect respondents' choices in the two multiple price list games during the follow-up survey. Because the information provided to respondents may impact MWTP differently for different groups, we only plot MWTP from the control group in this figure.

Figure A11: Respondent Misperceptions and Long-Run Changes in Air Quality



Notes: These figures test whether respondents' misperceptions about their relative air pollution exposure are explained by long-run changes in particulate matter levels. Panel (a) plots the absolute error in air pollution percentile beliefs against the 10-year change in air pollution levels in the respondent's ZIP code. Specifically, this long difference compares the 2004 to 2006 pollution levels against the 2014 to 2016 levels courtesy of [Meng, Li, Martin, Van Donkelaar, Hystad, and Brauer \(2019\)](#). Panel (b) uses the same variation in pollution, but plots residual errors in beliefs after controlling for respondents' current air pollution percentile, measured in 10-percentile bins. I.e., panel (b) tests whether 10-year changes in pollution explain misperceptions, controlling for current pollution levels. Panels (c) and (d) repeat this exercise, but use a 20-year difference: 1994 to 1996 versus 2014 to 2016.

Appendix Tables

Table A1: Balance Table, Race

	Control	Life Expectancy	Percentile	PM _{2.5} Level	N
<i>American Indian or Alaska Native</i>	0.160	0.360	0.320	0.160	25
<i>Asian</i>	0.266	0.223	0.244	0.266	349
<i>Black or African</i>	0.242	0.217	0.252	0.289	322
<i>Native Hawaiian or Other Pacific Islander</i>	0.000	0.429	0.286	0.286	7
<i>White</i>	0.254	0.253	0.247	0.246	2882
<i>2 or more races</i>	0.252	0.239	0.290	0.219	155
<i>Other</i>	0.188	0.250	0.219	0.344	64

Notes: This table reports the share of each race group in each of the four survey arms. The final column reports the total number of respondents in a given row's race group.

Table A2: Balance Table, Income

	Control	Life Expectancy	Percentile	PM _{2.5} Level	N
<i>Low</i>	0.254	0.244	0.256	0.246	1446
<i>Middle</i>	0.240	0.249	0.256	0.256	1397
<i>High</i>	0.266	0.252	0.228	0.254	961

Notes: This table reports the share of each income group in each of the four survey arms. We define high income as those with household incomes above \$100,000 per year; low-income respondents are those in households making below \$50,000 per year. The final column reports the total number of respondents in a given row's income group.

Table A3: Pollution Exposure by Race

Group	Mean PM _{2.5} ($\mu g/m^3$)
<i>All</i>	7.94
<i>White</i>	7.75
<i>Black</i>	8.24
<i>Asian</i>	8.63
<i>Latino</i>	8.44

Notes: This table displays the average particulate matter exposure by race for the 3,804 respondents who completed both the initial and follow-up surveys. PM_{2.5} measurements reflect the average of Di et al. (2016) and van Donkelaar et al. (2019). We use each individual's ZIP code centroid when converting the rasters from these two data sources to estimates of outdoor pollution exposure.

Table A4: Attrition

Statistic	Initial Survey	Follow-Up Survey	p-value
<i>N</i>	5137	3804	
<i>Share Asian</i>	0.091	0.093	0.816
<i>Share White</i>	0.707	0.727	0.080
<i>Share Black</i>	0.084	0.081	0.571
<i>Mean PM_{2.5}, White</i>	7.733	7.724	0.818
<i>Mean PM_{2.5}, Black</i>	8.277	8.265	0.888
<i>Mean PM_{2.5}, Asian</i>	8.535	8.603	0.578
<i>Share Control</i>	0.25	0.252	0.822
<i>Share LE Treatment</i>	0.251	0.248	0.776
<i>Share Percentile Treatment</i>	0.248	0.249	0.941
<i>Share PM_{2.5} Level Treatment</i>	0.252	0.252	0.989

Notes: This table compares the demographics, pollution exposure, and randomization of the 5,137 individuals who completed the initial survey and were eligible for the follow-up survey with the 3,804 individuals who successfully completed both the initial and follow-up surveys.

Table A5: Correlation in Respondent Errors

	<i>PM Error</i>	<i>Percentile Error</i>	<i>LE Error</i>
<i>PM Error</i>	1.000	0.011	-0.010
<i>Percentile Error</i>		1.000	-0.009
<i>LE Error</i>			1.000

Notes: This table reports the individual-level correlation coefficient in absolute errors across our three main elicitations: beliefs about PM_{2.5} levels, beliefs about pollution percentile, and beliefs about the gradient between pollution and life expectancy. When calculating errors, we take the “correct” answer to each question to be the mean of the two pieces of information that a given individual could have been shown during the information provision stage.

Table A6: Air Purifier WTP Summary Statistics

Group	Purifier A		Purifier B	
	Mean (\$)	SD (\$)	Mean (\$)	SD (\$)
<i>Full Sample</i>	33.08	28.12	55.27	39.58
<i>White</i>	32.95	27.96	56.19	40.33
<i>Black</i>	37.07	31.59	54.20	36.95
<i>Hispanic/Latino</i>	34.96	28.04	55.64	37.96
<i>Asian</i>	30.03	25.99	47.71	35.17
<i>High Income</i>	33.42	27.93	58.75	40.70
<i>Middle Income</i>	34.83	28.25	56.95	39.82
<i>Low Income</i>	31.17	28.01	51.35	38.26

Notes: This table shows summary statistics for participants’ willingness to pay for two air purifiers. Purifier A is a low-capacity air purifier, and purifier B is a high-capacity air purifier. The difference in the implied reduction in ambient air pollution depends on each respondent’s home size.

Appendix A Alternative Specifications and Robustness

Appendix A.1 Belief Updating and Source-Randomization

In the body of this paper, we instrument for posterior beliefs using the difference between an individual’s signal and prior, interacted with a treatment indicator. Our pre-analysis plan also described an additional and similar IV strategy where we would instead use the variation induced by the information *source*. Below we show that these first-stage relationships are generally weak, which motivates our use of the simpler signal vs. prior instrument.

Equation (10) tests whether posterior beliefs respond to the signal they are shown separately from the signal that they would have been shown in a different realization of the source-randomization:

$$x_{i,post}^k = \theta_0 + \theta_1^k z_i^k \underbrace{(s_{i,signal}^k - x_{i,prior}^k)}_{\text{signal}} + \theta_2^k z_i^k \underbrace{(s_{i,alt}^k - x_{i,prior}^k)}_{\text{alt signal}} + \theta_3^k (s_{i,signal}^k - x_{i,prior}^k) + \eta_i \quad (10)$$

The variable definitions are identical to equation (6), with the additional variable, $(s_{i,alt}^k - x_{i,prior}^k)$, representing the signal that treated individuals *would have been shown* in the other possible state of the source-randomization (recall that treated individuals were shown one of two possible information sources).

Table A7 reports the results of this exercise; the coefficient of interest in each column is $1[treated] \times (signal_i - prior_i)$. For each of the three beliefs, there is a positive coefficient on the interaction between treatment and the shown signal. Only for beliefs about pollution percentile, however, is this relationship strong enough to provide a valid instrument: The first-stage F-statistics for beliefs about pollution levels and life expectancy are 1.52 and 2.47, respectively. Given that the first-stage is much stronger for the simple Bayesian regressions (see Table 6), we do not use the source-randomization instrument in our main willingness to pay regressions.

Table A7: Belief Updating, Shown vs. Alternative Signal

	(1) Pollution Percentile	(2) Pollution Level	(3) Life Expectancy
<i>(Signal - Prior)</i>	-0.20* (0.02)	-0.23* (0.06)	-0.56* (0.02)
<i>Treated*(Signal - Prior)</i>	0.42* (0.04)	0.54 (0.44)	0.29 (0.18)
<i>Treated*(Alt - Prior)</i>	-0.07* (0.04)	-0.51 (0.43)	0.01 (0.18)
First-Stage F	97.91	1.52	2.47
N	1904	1914	1901

Notes: This table shows the results of applying equation (10) to beliefs about air pollution. The outcome variable in each column is the posterior belief for the corresponding question. Each regression includes respondents from two of the four survey arms: the arm that received information about the corresponding question, and the pure control group. *Signal* is the information that individuals either were shown or would have been shown had they been randomized into the treatment arm; *Alt* is the information that those randomized into the treatment were not shown (i.e., the information that a given respondent would have been shown in the other possible realization of the source-randomized information provided in the survey). Standard errors are heteroskedasticity-robust; * denotes results that are significant at the 0.05 level. Note that the significance levels in this table are not corrected for multiple hypotheses, as the results presented in this table are descriptive and do not directly test any of our main hypotheses.

Appendix A.2 Differences in MWTP for Clean Air Conditional on Income

Given the gap in income between Black and White households in the US and the fact that we document an income gradient in MWTP for clean air, it is natural to ask whether the difference in MWTP shown in column 1 of Table 5 is a consequence of differences in income across groups. In Table A8 below, we re-estimate the difference in MWTP between Black and White respondents while controlling for income. Specifically, we include fixed effects for each of the 16 self-reported income bins. Conditional on income, the Black-White gap in MWTP for clean air remains substantial — on average, it is roughly 20% of the mean MWTP for clean air.

Table A8: Differences in MWTP for Clean Air by Race, Conditional on Income

	MWTP	MWTP
1(<i>White</i>)	0.27* (0.08)	0.21* (0.08)
Income Controls		✓

Notes: Column 1 of this table reprints the results from Column 1 of Table 5, which is a simple bivariate regression of MWTP for clean air on an indicator for whether a respondent is White. Column 2 includes flexible controls for income — i.e., a fixed effect for each of the 16 income bins from our self-reported income variable. Both columns restrict the sample to only respondents who are either White or Black. Heteroskedasticity-robust standard errors are in parentheses; * denotes results that are significant at the 0.05 level. Note that the significance levels in this table are not corrected for multiple hypotheses, as the results presented in this table are descriptive and do not directly test any of our main hypotheses.

Appendix A.3 Robustness Checks for WTP Interpolation

We use a logit regression (equation Appendix A.3) to estimate WTP for each air purifier and each individual in our panel. For never-switchers — individuals who always prefer the air purifier to the gift card, or vice versa — WTP is not defined. For never-switchers who always prefer the gift cards, we assume a zero valuation; we see this as a weak assumption given that the lowest gift card value is \$1. We do not, however, have a natural upper bound for those who always prefer the air purifier. In the body of the paper, we assume that never-switchers have a WTP of \$150 (recall that the highest value of gift card was \$130). Below, we present robustness checks from two alternate assumptions about those who always prefer the air purifier: a “low” WTP of \$130 and a “high” WTP of \$200.

Table A9 reproduces our WTP summary statistics (see Table A6) for each of the two air purifiers under each of these alternatives. Table A11 and Table A10 reproduce our instrumental variables results from Table 6 for the “high” and “low” WTP interpolation values, respectively. Our results are robust to these alternate assumptions about the WTP of never-switchers: mean WTP is not sensitive to the value we use for interpolation, nor are our instrumental variables results.

Table A9: Air Purifier WTP Robustness

Group	Purifier A			Purifier B		
	(1) WTP A	(2) WTP A, \$130	(3) WTP A \$200	(4) WTP B	(5) WTP B, \$130	(6) WTP B, \$200
<i>Full Sample</i>	33.08	32.77	33.86	55.27	53.71	59.18
<i>White</i>	32.95	32.65	33.70	56.19	54.50	60.42
<i>Black</i>	37.07	36.44	38.62	54.20	52.83	57.61
<i>Hispanic/Latino</i>	34.96	34.70	35.61	55.64	54.34	58.88
<i>Asian</i>	30.03	29.86	30.47	47.71	46.91	49.72
<i>High Income</i>	33.42	33.13	34.15	58.75	56.98	63.18
<i>Middle Income</i>	34.83	34.52	35.58	56.95	55.27	61.14
<i>Low Income</i>	31.17	30.84	32.00	51.35	50.03	54.63

Notes: Column 1 and 4 of this table reproduce the mean WTP summary statistics from Table A6. Columns 2 and 3 provide summary statistics for air purifier A using \$130 and \$200 as interpolation values for never switchers, respectively. Columns 5 and 6 repeat this exercise for air purifier B.

Table A10: IV Results with a \$130 Interpolation Value for Never-Switchers

	(1) MWTP	(2) MWTP	(3) MWTP
<i>Percentile beliefs</i>	0.00 (0.01)		
<i>Pollution level beliefs</i>		0.20 (0.43)	
<i>Life expectancy beliefs</i>			0.24 (0.09)
First-Stage F	142.66	0.65	229.55
N	1904	1912	1901

Notes: This table reproduces the results from Table 6 using MWTP values that use a \$130 interpolation value for never-switchers who always chose the air purifier over the gift cards.

Table A11: IV Results with a \$200 Interpolation Value for Never-Switchers

	(1) MWTP	(2) MWTP	(3) MWTP
<i>Percentile beliefs</i>	0.00 (0.01)		
<i>Pollution level beliefs</i>		0.17 (0.40)	
<i>Life expectancy beliefs</i>			0.37 (0.14)
First-Stage F	142.66	0.65	229.55
N	1904	1912	1901

Notes: This table reproduces the results from Table 6 using MWTP values that use a \$200 interpolation value for never-switchers who always chose the air purifier over the gift cards.

Table A12: Beliefs and MWTP by Treatment Arm


	Control	Rank Treatment	PM Treatment	LE Treatment
<i>MWTP</i>	1.09 (1.55)	1.04 (1.42)	1.00 (1.42)	0.95 (1.24)
<i>Percentile Posterior</i>	53.50 (22.02)	45.11 (23.72)	54.49 (21.77)	51.21 (21.92)
<i>Percentile, Measured</i>	34.36 (23.69)	33.24 (23.35)	33.77 (23.52)	35.58 (23.67)
<i>PM_{2.5} Posterior</i>	11.40 (11.76)	15.19 (22.12)	11.81 (18.95)	11.96 (18.65)
<i>PM_{2.5}, Measured</i>	7.87 (1.58)	7.95 (1.59)	7.83 (1.59)	7.97 (1.58)
<i>LE Posterior</i>	1.96 (1.01)	1.98 (1.02)	2.00 (1.04)	1.35 (1.05)
<i>LE, Measured</i>	0.18 (0.09)	0.18 (0.09)	0.18 (0.09)	0.18 (0.09)

Notes: This table displays the mean for MWTP, posterior beliefs, and corresponding measured pollution values for each of the four survey randomization arms: the control group (column 1), the group that received information about their pollution percentile (column 2), the group that received information about their pollution level (column 3), and the group that received information about the relationship between air pollution and life expectancy (column 4). The standard deviation is listed in parentheses below each mean.

Appendix B Long-Run Follow-Up


Eight weeks after the initial survey we launched a second “long-run” follow-up survey. The goal of this survey was to gauge whether or not perceptions about air pollution would revert back to initial beliefs given a longer lag time. To that end, we again asked respondents to answer questions Q1 through Q3. Just as in the first follow-up, these questions were not incentivized. We plot the results of each of these questions below, along with the initial and first follow-up responses for reference.

Figure A12: 2 Month Follow-Up Beliefs About Relative Air Pollution Exposure




Notes: This figure plots respondents' beliefs about their relative air pollution exposure during the first follow-up survey and the long-run follow-up survey. Perceptions for the control group are plotted in red. Perceptions for the group that received information about the relative air quality in their ZIP code are plotted in blue for the 2-week follow-up, and purple for the 8-week follow-up. Perceptions about relative air quality for the two other treatment arms (air pollution levels and life expectancy) are not included in this figure. The dashed line is the locus where perceptions within an air pollution bin equal measured relative air pollution.

Figure A13: 2 Month Follow-Up Beliefs About Air Pollution Exposure



Notes: This figure plots respondents' beliefs about their air pollution exposure in levels during the first follow-up survey and the long-run follow-up survey. Perceptions for the control group are plotted in red. Perceptions for the group that received information about the air pollution level in their ZIP code are plotted in blue for the first follow-up, and in purple for the second follow-up. Perceptions about air quality for the two other treatment arms (air pollution levels and life expectancy) are not included in this figure. The dashed line is the locus where perceptions within an air pollution bin equal measured air pollution.

Figure A14: 2 Month Follow-Up Beliefs About the Health Impacts of Air Pollution



Notes: This figure plots respondents' perceptions about the relationship between air pollution and average life expectancy reported during the first follow-up survey and the long-run follow-up survey. Perceptions for the control group are plotted in red. Perceptions for the group that received information about the pollution-health gradient plotted in blue for the first follow-up, and in purple for the second follow-up. Perceptions about the pollution-health gradient for the two other treatment arms are not included in this figure.

Appendix C Additional Survey Details

This appendix provides examples of the key survey instrument questions and information screens. The formatting has been slightly adjusted from the original Qualtrics survey for readability. Duplicates of the initial and follow-up surveys are also available at the following links: [\[initial survey\]](#) [\[follow-up survey\]](#).

Appendix C.1 Initial Survey

Following a research information and informed consent landing page, respondents were asked to provide information about their home state and ZIP code.

What state do you live in?

[Drop-down menu]

What is your ZIP code?

[Text box]

We then tell respondents that the following questions will be incentivized:


The next few questions will ask you information about your ZIP code (94705).

It is in your best interest to answer these questions as accurately as possible – the 10% of respondents with the most accurate responses will receive a \$2 bonus.

You will not be able to go back and change your answers.

The next series of questions elicit priors for our three beliefs of interest – relative air pollution exposure, air pollution levels, and the average pollution-health gradient.

Question 1: *How would you rate **air pollution in your ZIP code** relative to **other ZIP codes within your state**?*



See [this website](#) if you would like a map of ZIP codes in your state.

Question 2: *The Clean Air Act regulates air quality in the United States. The current national air quality standard for particulate air pollution is 12 micrograms per cubic meter ($\mu\text{g}/\text{m}^3$). Areas with annual average air quality that exceed this threshold are subject to stringent pollution requirements.*

What do you think the average air pollution level in your ZIP code is?

[Short-answer text box that requires numerical answer]

To elicit beliefs about the connection between pollution and health, we first ask respondents whether they think the latter is related to the former:

If average air pollution in the US were 50% higher than it is right now, do you think this would impact US life expectancies?


[Yes, No]

If the respondent answers “No” in the above question, they are coded as having a belief that there is zero relationship between air pollution and life expectancy. If they respond “Yes” they are shown the following question:

Smoking¹, alcohol consumption², and motor vehicle accidents³ all decrease life expectancy.

If average air pollution in the US were 50% higher than it is right now, by how much do you think this would impact life expectancy?

Use the slider below to answer this question. The reduction in life expectancy from some other causes of death are included for reference.



Following these initial questions, respondents are randomized into one of the four survey arms. Three of these arms receive information corresponding to one of the three main survey questions (see examples below). The pure control arm proceeds directly to the closing information screen, also shown below. Each of the examples below that provide information about air pollution levels use the ZIP code 94705. Also recall that information was source-randomized, so conditional on their treatment arm, there were two possible pieces of information that could be shown to each individual.

Thank you for answering these questions.

*You will now randomly be shown the answer to **one** of the questions that you answered above.*

Please take a moment to read this information carefully.

Arm 1: Information about air pollution percentile. Respondents randomized into this arm receive information from either Di et al. (2016) or Van Donkelaar et al. (2019). Within each source, the information that is shown is a function of a respondent's ZIP code.

Di et al. (2016):

You answered that pollution in your ZIP code was better than 45% of ZIP codes within your state.

*According to satellite data¹, the particulate air pollution in your ZIP code is **better than 50%** of ZIP codes within your state.*

Van Donkelaar et al. (2019):

You answered that pollution in your ZIP code was better than 45% of ZIP codes within your state.

*According to satellite data¹, the particulate air pollution in your ZIP code is **better than 37%** of ZIP codes within your state.*

Arm 2: Information about air pollution levels. Respondents randomized into this arm receive information from either Di et al. (2016) or Van Donkelaar et al. (2019). Within each source, the information that is shown is a function of a respondent's ZIP code.

Di et al. (2016):

You answered that the average particulate air pollution level in your ZIP code is $12\mu\text{g}/\text{m}^3$.

According to satellite data¹, the average particulate air pollution level in your ZIP code is $7.79\ \mu\text{g}/\text{m}^3$.

Van Donkelaar et al. (2019):

You answered that the average particulate air pollution level in your ZIP code is $12\mu\text{g}/\text{m}^3$.

According to satellite data¹, the average particulate air pollution level in your ZIP code is $7.58\ \mu\text{g}/\text{m}^3$.

Arm 3: Information about life expectancy. Respondents randomized into this arm receive information from either Krewski et al. (2009) or Anderson (2020). This information is not a function of a respondent's ZIP code.

Anderson (2020):

*You answered that **increasing air pollution in the US by 50% would shorten life expectancies by 1.5 years.***

This is an overestimate

*According to academic research¹, increasing air pollution in the US by 50% would shorten life expectancies by **0.09 years.***

Krewski et al. (2009):

*You answered that **increasing air pollution in the US by 50% would shorten life expectancies by 1.5 years.***

This is an overestimate

*According to academic research¹, increasing air pollution in the US by 50% would shorten life expectancies by **0.27 years.***

Closing screen:

We thank you for your time spent taking this survey.

Your response has been recorded; your completion code is [8-character Prolific code].

We will contact you about a follow-up survey in 2-4 weeks.

Appendix C.2 Follow-up survey

The follow-up survey begins by re-asking the three main belief questions from the initial survey to measure updating. After eliciting these follow-up beliefs, we asked respondents a question about their home size in order to calibrate the information that we show them about the effectiveness of two air purifiers. Respondents then complete the real-stakes multiple price list game. To conclude, we collect demographic data on respondents.

How many square feet is your home?

[Text box requiring numeric answer]

The following question will show you some information about air purifiers, and then ask you whether you would rather receive an air purifier or an Amazon gift card of differing values.

It is in your best interest to answer these questions truthfully: 1 in 50 respondents will actually receive one of the items that they choose.

The information provided to each respondent depends on the size of a respondent's home. We assume 9-foot ceilings to convert square footage to volume, and then use each air purifier's clean air delivery rate (CADR) to estimate the steady-state reduction in indoor air pollution under the assumption of 1 air change per hour. Someone with a 1000 square foot living space would receive the following information:

*An air purifier's effectiveness depends on the **capacity of the air purifier** and the **size of the home** that it is in. In the following questions, you will be asked to make decisions about these two types of air purifiers:*

Air Purifier A

Based on your home size, this air purifier would reduce indoor air pollution in your home by 22%



Air Purifier B

Based on your home size, this air purifier would reduce indoor air pollution in your home by 49%



Respondents then complete both of the following multiple price lists. The order of the items shown, the direction of gift card values (ascending or descending), and the order of columns within each list (gift cards first or air purifiers first) is randomized across respondents.

*There are 10 pairs of items below. For each case, choose the item that you would prefer to receive.
Recall that **Air Purifier B** would reduce indoor air pollution in your home by **49%***

Choice 1	Air Purifier B	\$1 Amazon Gift Card
Choice 2	Air Purifier B	\$5 Amazon Gift Card
Choice 3	Air Purifier B	\$15 Amazon Gift Card
Choice 4	Air Purifier B	\$25 Amazon Gift Card
Choice 5	Air Purifier B	\$35 Amazon Gift Card
Choice 6	Air Purifier B	\$45 Amazon Gift Card
Choice 7	Air Purifier B	\$60 Amazon Gift Card
Choice 8	Air Purifier B	\$75 Amazon Gift Card
Choice 9	Air Purifier B	\$95 Amazon Gift Card
Choice 10	Air Purifier B	\$130 Amazon Gift Card

*There are 10 pairs of items below. For each case, choose the item that you would prefer to receive.
Recall that **Air Purifier A** would reduce indoor air pollution in your home by **22%***

Choice 1	Air Purifier A	\$1 Amazon Gift Card
Choice 2	Air Purifier A	\$5 Amazon Gift Card
Choice 3	Air Purifier A	\$15 Amazon Gift Card
Choice 4	Air Purifier A	\$25 Amazon Gift Card
Choice 5	Air Purifier A	\$35 Amazon Gift Card
Choice 6	Air Purifier A	\$45 Amazon Gift Card
Choice 7	Air Purifier A	\$60 Amazon Gift Card
Choice 8	Air Purifier A	\$75 Amazon Gift Card
Choice 9	Air Purifier A	\$95 Amazon Gift Card
Choice 10	Air Purifier A	\$130 Amazon Gift Card

Finally, we ask respondents to answer a series of demographic questions. These questions supplement demographic data on respondents provided by Prolific.

How old are you?

[Text box requiring a numerical answer]

What is your race?

- ☐ White
- ☐ Black
- ☐ American Indian or Alaska Native
- ☐ Asian
- ☐ Native Hawaiian or Other Pacific Islander
- ☐ 2 or more races
- ☐ other

Are you of Hispanic, Latino, or Spanish origin?

[Yes, No]

What is your gender identity?

- ☐ Male
- ☐ Female
- ☐ Nonbinary
- ☐ Prefer not to say

What is your approximate household income?

- ☐ \$0 to \$9,999
- ☐ \$10,000 to \$19,999
- ☐ \$20,000 to \$29,999
- ☐ \$30,000 to \$39,999
- ☐ \$40,000 to \$49,999
- ☐ \$50,000 to \$59,999
- ☐ \$60,000 to \$69,999
- ☐ \$70,000 to \$79,999
- ☐ \$80,000 to \$89,999
- ☐ \$90,000 to \$99,999
- ☐ \$100,000 to \$124,999
- ☐ \$125,000 to \$149,999
- ☐ \$150,000 to \$174,999
- ☐ \$175,000 to \$199,999
- ☐ \$200,000 to \$249,999
- ☐ Over \$250,000

Do you already own an air purifier?

[Yes, no]

Appendix D Pre-Analysis Plan

The pre-analysis plan for this project was registered with the AEA RCT Registry on January 12th, 2023 (*AEARCTR-0010435, The Role of Information in the Willingness to Pay for Clean Air*).

There are several minor ways in which the analysis in this paper differs from what we outlined in the pre-analysis plan. First, because participants update their beliefs in response to information treatment, but not in such a way that the source-randomization predicts posterior beliefs, one set of instruments described in Section 4 of the pre-analysis plan is not valid. Instead, we use the difference between the signal and the prior as an instrument for follow-up beliefs (as described in Section 5.3). Second, our pre-analysis plan describes using a single IV regression where we simultaneously test whether changes in beliefs lead to changes in individuals’ marginal willingness to pay for clean air. Given that identification in 2SLS with multiple treatment arms requires strong assumptions (see Bhuller and Sigstad, 2022), we instead separately estimate three single-variable IV regressions to answer the questions outlined in Hypothesis Group 2. Lastly, we used *Prolific* rather than *Mechanical Turk* to conduct this study because of complications with a new Amazon billing system that was launched in late 2022.

Appendix D.1 Hypotheses

We have three main research questions: (i) What do people know about their air pollution exposure and its health impacts?; (ii) Does the accuracy of beliefs about air pollution differ across race or income groups?; and, (iii) How do changes in beliefs about air pollution impact demand for clean air? We map these research questions into a series of pre-specified hypotheses posted on the AEA’s RCT registry, *AEARCTR-0010435*.

Prespecification presents a tradeoff for researchers: adding additional tests allows for additional credible discoveries, but can reduce the power of the other hypotheses a researcher wishes to conduct due to adjustments for multiple-hypothesis testing. In our context, we have several outcomes – measures of misperception, belief updating, and willingness to pay for clean air – and several race and income subgroups where we could test for differences in these outcomes. To balance this tradeoff, we follow the procedure described by Anderson and Magruder (2017) and designate “main” hypotheses, which speak directly to the research questions described above, and “auxiliary” hypotheses which answer other natural questions that come up in our setting. This procedure preserves statistical power for testing the main hypotheses and then uses a split-sample technique to efficiently test the auxiliary hypotheses. The following sections list each hypothesis and details how we adjust our inference thresholds to account for multiple tests.

Appendix D.1.1 Main Hypotheses

We test 3 main groups of hypotheses, some of which contain multiple comparisons, which in total lead to 11 primary tests. We demarcate the tables where these hypotheses are tested with a dagger (†).

Hypothesis Group 1: *The accuracy of beliefs about air pollution is the same across race groups and income groups.* For each subgroup and pollution question, we measure accuracy using the mean absolute error, where we take the “correct” answer to each question to be the mean of the two pieces of information that a given individual could have been shown during the information provision stage. The corresponding null hypothesis for two groups j and l is that the mean absolute error in misperceptions is equal for these two groups. We test for differences in perceptions for each of the survey’s three questions (relative pollution exposure, pollution exposure in levels, and life expectancy) separately. Comparing these means therefore yields 6 total hypotheses (three pairwise comparisons across income groups; three pairwise comparisons across race groups). We evaluate this comparison using the (multiple-hypothesis-corrected) p -value from a simple bivariate regression of mean absolute error on a group indicator.

Hypothesis Group 2: *Beliefs about air pollution (exposure, relative exposure, and health effects) affect marginal willingness to pay for reductions in air pollution.* For any belief type, the corresponding null hypothesis is that the coefficient on posterior beliefs in Equation (9) is zero. Equation (9) is implemented using two stage least squares where the information provision is used as an instrument for posterior beliefs. This hypothesis group consists of three tests (i.e., one for each belief), evaluated using the multiple-hypothesis-corrected p -value on the coefficient of interest in each regression.

Hypothesis Group 3: *The marginal willingness to pay for air pollution reductions is the same across race and income groups.* This group consists of 2 tests. Each test uses marginal willingness to pay estimated via equation 8 as the left-hand side variable and a group indicator — either $I(\text{White})$ or $I(\text{low} - \text{income})$ — as

the right-hand side variable. We evaluate a given hypothesis using the corrected p -value associated with the corresponding indicator variable.

Appendix D.2 Auxiliary Hypotheses

To avoid diluting the power of our main hypotheses tests, we use a split-sample approach to screen out auxiliary hypotheses that may be underpowered to be able to reject a null. Specifically, we first test all auxiliary hypotheses in a 35% “exploratory” sample. These results from these exploratory sample tests are reported in Tables A13 through A16, below.

For the tests that passed the screening threshold of significance ($t = 1.6$) in the exploratory sample, we then confirm or reject these hypotheses using the 65% hold-out sample. As suggested by Anderson and Magruder (2017), we filed a separate pre-analysis plan for these auxiliary hypotheses (*AEARCTR-0011970*). Below we list the 6 hypotheses that passed the threshold for significance in the exploratory sample. Note that we also specify a direction for each test because we use a one-sided t-test to confirm/reject each hypothesis in the 65% hold-out sample.²⁸

1. **Auxiliary Hypothesis 1:** The mean WTP for air purifier B is the same for high- and low-income respondents. **Direction:** In the exploratory sample, the WTP for air purifier B was higher for high-income respondents ($t \geq 1.6$).
2. **Auxiliary Hypothesis 2:** The mean marginal willingness to pay for clean air (i.e., a 1% reduction in household air pollution exposure) is the same for White and Asian respondents. **Direction:** In the exploratory sample, the MWTP for clean air was higher for White respondents ($t \geq 1.6$).
3. **Auxiliary Hypothesis 3:** The mean accuracy of beliefs about the pollution-health gradient is the same for White respondents and Asian respondents. **Direction:** In the exploratory sample, the absolute error in beliefs was higher for White respondents ($t \geq 1.6$).
4. **Auxiliary Hypothesis 4:** The mean WTP for air purifier B is the same for White and Asian respondents. **Direction:** In the exploratory sample, the WTP for air purifier B was higher for White respondents ($t \geq 1.6$).
5. **Auxiliary Hypothesis 5:** The mean accuracy of beliefs about particulate matter exposure (in levels) is the same for White respondents and Hispanic/Latino respondents. **Direction:** In the exploratory sample, the absolute error in beliefs was higher for Hispanic/Latino respondents ($t \geq 1.6$).
6. **Auxiliary Hypothesis 6:** Black and White respondents exhibit the same belief updating behavior (i.e., the coefficient on the relationship between one’s updating and the signal that they were shown) with respect to beliefs about the pollution-health gradient. **Direction:** In the exploratory sample, White respondents updated their beliefs more than did Black respondents ($t \geq 1.6$).

Auxiliary hypotheses tested on the 35% exploratory sample:

²⁸Anderson and Magruder (2017) suggest using a one-sided test to improve power in the hold-out tests, and note that “the chance of finding a highly significant coefficient that goes in the opposite direction in the confirmation sample is extremely small.”

Table A13: Differences in WTP Across Race and Income Groups, Exploratory Sample

	(1) WTP	(2) WTP	(3) WTP	(4) WTP
<i>1(Hispanic)</i>	−2.67 (3.22)			
<i>1(Asian)</i>		−11.47 (3.09)		
<i>1(Black)</i>			−1.71 (3.84)	
<i>1(Low Income)</i>				−6.27 (2.81)
<i>t-stat</i>	−0.83	−3.71	−0.45	−2.24

Notes: This table uses the 35% exploratory sample to compare WTP across race and income groups. Columns (1) through (3) compare the WTP for air purifier “B” in the indicated race group to the WTP among White respondents. The final column compares low- and high-income respondents. The final row of this table reports *t*-statistics used to determine whether observed differences in the exploratory sample are large enough to justify testing a given hypothesis in the hold-out sample.

Table A14: Differences in Updating Across Race Groups, Exploratory Sample

	(1) Percentile	(2) PM	(3) LE	(4) Percentile	(5) PM	(6) LE
<i>1(Hispanic)</i>	2.15 (1.80)	6.91 (2.56)	0.01 (0.09)			
<i>1(Asian)</i>				2.24 (1.99)	1.25 (1.92)	−0.22 (0.08)
<i>t-stat</i>	1.20	2.70	0.11	1.13	0.65	−2.65

Notes: This table uses the 35% exploratory sample to compare belief updating between Hispanic and White respondents and between Asian and White respondents. The final row of this table reports *t*-statistics used to determine whether observed differences in the exploratory sample are large enough to justify testing a given hypothesis in the hold-out sample.

Table A15: Differences in MWTP Across Race Groups, Exploratory Sample

	(1) MWTP	(2) MWTP
<i>1(Hispanic)</i>	−0.12 (0.10)	
<i>1(Asian)</i>		−0.27 (0.10)
<i>t-stat</i>	−1.24	−2.67

Notes: This table uses the 35% exploratory sample to compare MWTP across race and income groups (also see Table 5). In each regression, the reference group is White respondents. The final row of this table reports *t*-statistics used to determine whether observed differences in the exploratory sample are large enough to justify testing a given hypothesis in the hold-out sample.

Table A16: Differences in Updating Across Race and Income Groups, Exploratory Sample

	(1) Percentile	(2) PM	(3) LE	(4) Percentile	(5) PM	(6) LE
$1(White)*Treat*(Signal - Prior)$	-0.03 (0.18)	0.10 (1.27)	0.19 (0.11)			
$1(Low-Income)*Treat*(Signal - Prior)$				-0.10 (0.14)	-0.42 (0.45)	-0.01 (0.09)
t -stat	-0.17	0.08	1.70	-0.69	-0.94	-0.13

Notes: Columns (1) through (3) of this table compare differences in belief updating between White and Black respondents; columns (4) through (6) compare differences in belief updating between low- and high-income respondents. The final row of this table reports t -statistics used to determine whether observed differences in the exploratory sample are large enough to justify testing a given hypothesis in the hold-out sample.

Auxiliary hypotheses tested on the 65% hold-out sample:

Table A17: Differences in WTP Across Race Groups, Hold-out Sample

	WTP	WTP
$1(Asian)$	-6.88* (2.66)	
$1(Low\ Income)$		-7.97* (2.05)
t -stat	-2.58	-3.88

Notes: This table uses the 65% hold-out sample to compare WTP for air purifier “B” between White and Asian respondents.

Table A18: Differences in MWTP Across Race Groups, Hold-out Sample

	MWTP
$1(Asian)$	-0.24* (0.10)
t -stat	-2.45

Notes: This table uses the 65% hold-out sample to compare White and Asian respondents’ MWTP for reductions in air pollution.

Table A19: Differences in Belief Accuracy Across Race Groups, Hold-out Sample

	PM	LE
$1(Hispanic)$	-70.54 (72.18)	
$1(Asian)$		-0.10 (0.07)
t -stat	-0.98	-1.41

Notes: This table uses the 65% hold-out sample to test for the differences in the accuracy of initial beliefs across race groups observed in the exploratory sample. The first column compares accuracy of beliefs about particulate matter levels between White and Hispanic/Latino respondents. The second column compares the accuracy of beliefs about the relationship between air pollution and life expectancy between White and Asian respondents.

Table A20: Differences in Updating Across Race and Income Groups, Hold-out Sample

	LE
$1(White)*Treat*(Signal - Prior)$	0.08 (0.11)
t -stat	0.71

Notes: This table uses the 65% hold-out sample to test for the differences in belief updating detected in the exploratory sample.

Appendix D.3 Correcting for multiple hypotheses

We have 17 total hypotheses: 11 main hypotheses, and 6 auxiliary hypotheses that passed the screening threshold for significance ($t = 1.6$) in the exploratory sample. We correct for testing these hypotheses using the step-down procedure described by [Benjamini and Hochberg \(1995\)](#).²⁹ This process involves (1) ranking the m main hypotheses by their uncorrected p -values, (2) finding the hypothesis with the highest rank, j , for which the p -value, p_j is less than or equal to $(j/m) * (0.05)$, and (3) rejecting all hypotheses with rank less than or equal to j . In all tables where we report the results from one of these 17 tests, we present both the unadjusted and adjusted p -values. The results of the Benjamini-Hochberg step-down procedure are in Table [A21](#).

Table A21: Benjamini-Hochberg Multiple Hypothesis Corrections

Rank	Raw p-value	Hypothesis	Description	Table	Adjusted threshold	Significant
1	<0.001	Main 8	MWTP, Low vs. High Income	Table 5	0.003	Yes
2	<0.001	Auxiliary 2	WTP, High vs. Low Income	Table A17	0.006	Yes
3	<0.001	Main 1	Percentile Beliefs, White vs. Black	Table 3	0.009	Yes
4	<0.001	Main 2	PM Beliefs, White vs. Black	Table 3	0.012	Yes
5	0.001	Main 7	MWTP, White vs. Black	Table 5	0.015	Yes
6	0.007	Main 11	LE Beliefs and MWTP (IV)	Table 6	0.018	Yes
7	0.01	Auxiliary 1	WTP, White vs. Asian	Table A17	0.021	Yes
8	0.014	Auxiliary 3	MWTP, White vs. Asian	Table A18	0.024	Yes
9	0.037	Main 3	LE Beliefs, White vs. Black	Table 3	0.026	No
10	0.091	Main 6	LE Beliefs, Low vs. High Income	Table 3	0.029	No
11	0.158	Auxiliary 5	LE Beliefs, White vs. Asian	Table A19	0.032	No
12	0.329	Auxiliary 4	PM Beliefs, White vs. Hispanic/Latino	Table A19	0.035	No
13	0.401	Main 4	Percentile Beliefs, Low vs. High Income	Table 3	0.038	No
14	0.496	Auxiliary 6	LE Updating, White vs. Black	Table A20	0.041	No
15	0.51	Main 9	Percentile Beliefs and MWTP (IV)	Table 6	0.044	No
16	0.657	Main 10	PM Beliefs and MWTP (IV)	Table 6	0.047	No
17	0.672	Main 5	PM Beliefs, Low vs. High Income	Table 3	0.050	No

Notes: This table applies the Benjamini-Hochberg step-down procedure to 17 p -values from our main and auxiliary hypotheses. The “adjusted threshold” is the cutoff for significance that accounts for multiple hypotheses. It is equal to $(j/m) * 0.05$, where j is the rank of a given test, and m is the total tests. If the raw p -value is lower than this adjusted threshold, the test is significant at the 5% level, after accounting for multiple hypotheses.

²⁹[Benjamini and Yekutieli \(2001\)](#) show that this procedure controls for the false discovery rate under positively-correlated test statistics and claim that this is a plausible ex-ante assumption for the dependence of tests in multi-arm studies.