

NBER WORKING PAPER SERIES

HOUSEHOLD RESPONSES TO PRIVATE RISK INFORMATION

Benjamin Krebs
Matthew J. Neidell

Working Paper 34875
<http://www.nber.org/papers/w34875>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2026

For comments and suggestions, we thank Ludovica Gasse, Sefi Roth, Simon Luechinger, the Public Economics Group at the University of Basel, as well as all participants at the Workshop on Applied Environmental and Health Economics at Università Cattolica del Sacro Cuore, Rome (July 2025), and the research seminar of the University of Basel Faculty of Business and Economics (December 2025). Benjamin Krebs acknowledges funding by the Swiss National Science Foundation Grant Number PZ00P1_223786. The authors report no financial or other relationships that could be construed as a potential conflict of interest related to this research. The authors received no external financial or in-kind support for this work beyond what is disclosed in the paper. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Benjamin Krebs and Matthew J. Neidell. 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.

Household Responses to Private Risk Information
Benjamin Krebs and Matthew J. Neidell
NBER Working Paper No. 34875
February 2026
JEL No. D81, D83, Q53, Q55

ABSTRACT

We study how private information affects household responses to environmental risk. Using data from residential air quality monitors, we exploit the timing of monitor installation and high-frequency fine particulate matter (PM_{2.5}) readings to identify responses to new information about indoor pollution risk. We find that indoor PM_{2.5} concentrations decline by 2.5 ug/m³ over the 12 weeks following installation, conditional on contemporaneous outdoor pollution, with effects significantly larger among households with high initial indoor pollution. The indoor–outdoor pollution gradient declines over time, indicating that households become increasingly effective at mitigating exposure when marginal health damages are highest. Using machine learning techniques to infer cooking activity and air purifier adoption, we show that households respond primarily through durable defensive investments rather than reductions in pollution-generating behavior, with back-of-the-envelope calculations implying positive net benefits. Our results suggest that personalized risk information increases the salience of indoor pollution as a controllable risk for households, in contrast to spatially coarse public information that frames pollution primarily as an outdoor threat requiring avoidance.

Benjamin Krebs
University of Basel
Faculty of Business and Economics
benjamin.krebs@unibas.ch

Matthew J. Neidell
Columbia University
Department of Health Policy & Management
and NBER
mn2191@columbia.edu

1 Introduction

Economic decisions often depend on noisy or incomplete information about underlying risk. In many settings, publicly available signals are spatially coarse, while realized conditions—and therefore expected outcomes—vary at finer scales. For example, city-level crime statistics obscure safety conditions on specific streets, aggregate case counts of infectious disease fail to capture person-to-person transmission risk, and citywide heat warnings mask temperature variation across neighborhoods. As technological advances increasingly provide personalized information, an important question is how individuals update beliefs and adjust behavior when they receive signals that more closely align with their own exposure.

In this paper, we study how the arrival of private risk information shapes household behavior in the context of air pollution. The limitations of spatially aggregated information are particularly pervasive for environmental risks. Public information about air pollution has expanded substantially over the past two decades, with numerous studies showing that people respond to such public information by engaging in avoidance behavior (Neidell 2009; Ito and Zhang 2020; Barwick et al. 2024).¹ Outdoor measurements, however, may poorly reflect indoor exposures where Americans spend approximately 87% of their time (Klepeis et al. 2001). Although outdoor pollution enters indoor environments, this infiltration is incomplete, and additional pollution is generated indoors. Moreover, outdoor air pollution varies considerably even within small geographic areas, such that measurements from centrally located monitors may not accurately reflect personal exposures. As a result, individuals relying on public information face uncertainty about their true exposure.

Recent advances in low-cost sensor technology have begun to change information available to individuals. Residential air quality monitors allow households to observe pollution levels inside their home in real time, generating private, household-level measures of risk. Over the last decade, numerous U.S. households have installed these devices to track their indoor air quality. We obtain data from PurpleAir, the dominant platform for residential air quality monitoring in the U.S., to examine how households respond to private risk information.

Our empirical strategy exploits the timing of indoor air quality monitor installation as the arrival of new information about indoor pollution risk, allowing us to estimate subsequent changes in measured indoor air quality. Since households need time to process information and adopt mitigation strategies, this provides a natural baseline for identifying responses. Interpreting the resulting changes as causal responses to information requires that households had limited prior knowledge of their IAQ before monitor installation, a plausible assumption given that $\text{PM}_{2.5}$ is largely undetectable by human senses and households purchase monitors specifically to learn about IAQ. Using high-frequency measurements of $\text{PM}_{2.5}$, we track changes in indoor air quality from the moment households begin receiving information about their exposure. Furthermore, we flexibly

¹In the United States alone, the Environmental Protection Agency maintains a national network of over 5,000 active air quality monitors that provide real-time data on outdoor pollution concentrations. This information is widely disseminated through weather apps, news broadcasts, and dedicated air quality websites, reaching millions of Americans daily.

control for contemporaneous outdoor pollution to separate responses to information from mean reversion due to monitor adoption coinciding with major pollution episodes, such as wildfires.

We find that indoor $\text{PM}_{2.5}$ concentrations significantly decline following monitor installation, with average reductions of approximately $2.5 \mu\text{g}/\text{m}^3$ by the twelfth week after installation. This average effect, however, hides important heterogeneity across households. Improvements are concentrated among households with high initial indoor pollution, particularly when it exceeds outdoor pollution at the time information arrives. For example, among households facing initial outdoor pollution $\geq 12 \mu\text{g}/\text{m}^3$, those with indoor $\text{PM}_{2.5}$ levels exceeding outdoor levels experience declines by $22 \mu\text{g}/\text{m}^3$ after twelve weeks. Moreover, the indoor–outdoor pollution gradient declines over time, particularly during high outdoor pollution episodes, indicating that households become increasingly effective at limiting exposure when the marginal benefits of mitigation are greatest. This heterogeneity is consistent with active household mitigation rather than passive regression to the mean.

To understand the mechanisms underlying these responses, we use machine learning techniques to infer two key household behaviors: cooking episodes and air purifier adoption. We use outlier detection methods to identify cooking episodes, which generate distinctive spikes in indoor $\text{PM}_{2.5}$ uncorrelated with outdoor conditions. We employ k-means clustering on the prevalence of very low indoor pollution readings to identify households that adopt air purifiers.² Importantly, we exclude cooking episodes from the air purifier detection analysis since cooking-related spikes would obscure the baseline filtration effect of purifiers.

We find little evidence that households reduce pollution-generating activities such as cooking or meaningfully alter ventilation behavior over time. Instead, households respond primarily through durable defensive investments, particularly the adoption of air purifiers, which deliver persistent reductions in indoor pollution and are especially effective during periods of elevated outdoor pollution. When outdoor $\text{PM}_{2.5}$ is between $5\text{-}10 \mu\text{g}/\text{m}^3$, purifiers reduce indoor concentrations by approximately $2 \mu\text{g}/\text{m}^3$. This effect increases monotonically with outdoor pollution, reaching $11 \mu\text{g}/\text{m}^3$ when outdoor levels are $30\text{-}35 \mu\text{g}/\text{m}^3$ and $25 \mu\text{g}/\text{m}^3$ when outdoor pollution exceeds $35 \mu\text{g}/\text{m}^3$. These effects persist even during cooking episodes, providing additional validation of our classification approach, given that cooking episodes are excluded from our purifier detection algorithm. Back-of-the-envelope calculations suggest net benefits from air purifier adoption of $\$325\text{-}575$ in the first year alone.

Our results point to a key distinction between private and public risk information. While public monitoring networks provide widespread information about outdoor pollution, this portrays pollution as an external threat to avoid. Private indoor monitors change that perspective to an internal threat that is controllable through individual actions. This distinction is particularly important when spatially aggregated public information masks heterogeneity in household-level exposure.

²“Air purifier” refers to any filtration technology that removes $\text{PM}_{2.5}$ from indoor air, including portable air purifiers and central HVAC systems with high-efficiency filters (HEPA or MERV 13+). Note that we use “adoption” interchangeably with “usage” as we cannot identify them separately.

Our study contributes to three main strands of literature. First, our finding that personalized monitoring increases the salience of indoor pollution as a controllable risk parallels other salience interventions: tax salience affects consumer choices (Chetty et al. 2009), visual cost displays in health insurance influence plan selection (Handel 2013), smart electricity meters reduce energy consumption (Allcott 2011), and fitness trackers increase exercise (Royer et al. 2015). In each case, the intervention delivers information in a form that is easier for individuals to act upon. In our setting, private monitoring similarly makes pollution risk actionable by shifting household responses from avoidance toward mitigation.

Second, we contribute to the literature on behavioral responses to environmental risk. Much of the previous literature has examined temporary behavioral responses, such as changes in outdoor activities (Graff Zivin and Neidell 2009; Neidell 2009, 2010; Noonan 2014; Ward and Beatty 2016; Saberian et al. 2017; Barwick et al. 2024; Anderson et al. 2022; Shi et al. 2023; Fan 2024; Aguilar-Gomez 2025), which require ongoing effort and provide protection only during avoidance. We show that information provision triggers durable technology adoption, which provides continuous protection with minimal ongoing cost. While Ito and Zhang (2020) find that purifier purchases increase with ambient air pollution levels, our contribution is distinct: we show that information itself, independent of changes in actual pollution, triggers adoption. In our setting, outdoor pollution levels remain constant; only households’ knowledge of their exposure changes.

Our third contribution relates to the work on responses to information by examining private rather than public information. While previous research has largely focused on publicly provided information, such as air quality alerts, we show that households actively seek and respond to personalized information even in settings where public information is available. This reveals important limitations of public information systems: they rely on spatially coarse measures that often poorly reflect household-specific exposure. Our findings demonstrate that households recognize the value of personalized information, which motivates private investment.

Several recent studies examine household responses to private IAQ information. Greenstone et al. (2021) conducted a field experiment in Delhi, India, randomly assigning pollution monitors to over 4,000 households and offering subsidized air purifier rentals. They find that monitors did not affect air purifier take-up, though the authors note significant attrition and caution that results should be interpreted as suggestive. Sater et al. (2022) installed IAQ monitors in 281 French homes with wood-burning fireplaces and provided weekly static feedback on PM_{2.5} levels, finding that personalized information reduced indoor pollution by over 20%. Metcalfe and Roth (2025) conducted a randomized controlled trial of 258 households in London, where the control group received monitors with blinded displays while the treatment group saw real-time information. They find that real-time feedback reduced indoor PM_{2.5} by 17% overall and 34% during occupancy hours over a four-week period, primarily through increased ventilation via window opening.

We extend these findings in two ways. First, our observational approach leverages data from over 4.4 million observations from 2,249 households in California who voluntarily purchased PurpleAir monitors, following them for twelve weeks. While Metcalfe and Roth (2025) find that window

opening, a behavioral change requiring ongoing effort, is the primary mechanism over four weeks, we identify air purifier adoption as the dominant response over twelve weeks. The observation window matters: shorter-term studies may capture transient behavioral adjustments like ventilation, while longer-term observation reveals households’ willingness to invest in durable technology that provides continuous protection.

Second, our results suggest that both selection into monitor adoption and contextual factors appear important for understanding household responses. Unlike [Greenstone et al. \(2021\)](#), who find no effect on purifier take-up in a randomized sample in Delhi, India, we find strong evidence that information triggers defensive investments among households who voluntarily adopt monitors in California. This difference could reflect selection: households willing to invest in monitoring may respond more to the information, consistent with revealed preference for IAQ improvements. Moreover, contextual differences between the U.S. and India are also likely important, such as differences in the affordability of air purifiers. Taken together, these differences suggest important implications for the external validity of our findings, particularly given endogenous information acquisition.

The remainder of the paper proceeds as follows. [Section 2](#) provides background on IAQ and PurpleAir monitors and describes our data. [Section 3](#) presents our empirical strategy. [Section 4](#) reports our main results on changes in IAQ following monitor installation. [Section 5](#) examines the specific mitigation behaviors that households undertake. [Section 6](#) concludes.

2 Background and Data

2.1 Indoor Air Quality and Health

Americans spend approximately 87% of their time indoors, making IAQ a primary determinant of pollution exposure ([Klepeis et al. 2001](#)). Indoor $PM_{2.5}$ concentrations reflect both the infiltration of outdoor pollution and indoor emission sources, including cooking, cleaning, candles, and cigarette smoking. The relative importance of indoor and outdoor sources varies substantially across homes depending on building characteristics, ventilation practices, and occupant behavior, though the correlation is typically high ([Chen and Zhao 2011](#); [Krebs et al. 2021](#)). However, the relationship is far from one-to-one, and indoor levels can diverge substantially from ambient conditions. Hence, information can reduce personal exposure even when outdoor pollution remains unchanged.

Exposure to $PM_{2.5}$ has well-documented health effects. Short-term exposure increases respiratory and cardiovascular morbidity as well as many non-health endpoints ([Aguilar-Gomez et al. 2022](#)), while long-term exposure contributes to chronic disease and reduces life expectancy ([Pope III et al. 2009](#); [Chen et al. 2013](#); [Anderson 2020](#); [Barreca et al. 2021](#); [Deryugina and Reif 2023](#)). Recent evidence suggests health effects even at concentrations below regulatory standards ([Di et al. 2017](#)). Given these health impacts and the dominance of indoor exposure, understanding how households manage IAQ has the potential to produce substantial public health benefits.

2.2 PurpleAir Monitors

PurpleAir monitors emerged in 2015 as a low-cost device for measuring particulate matter concentrations. The monitors use laser particle counters to measure $\text{PM}_{2.5}$ and report readings every two minutes to a central database accessible through a public API. At approximately \$250 per unit, they cost a fraction of regulatory-grade monitors while maintaining reasonable accuracy (correlation of 0.88 with federal reference monitors; Kelly et al. 2017). We apply a set of conventional quality checks to ensure data reliability (see below).

The PurpleAir network has grown rapidly, particularly in regions affected by wildfires. California alone has over 10,000 monitors as of 2025, with substantial numbers in other U.S. states. Initially, most monitors were installed outdoors, though an increasing share are now placed indoors as households seek to understand their home air quality. These indoor monitors provide our primary data source.

Figure 1 illustrates the relationship between pollution episodes and monitor adoption during our study period from 2019 through 2021. Panel A shows that Google searches for “air purifier” (gray line) and “purple air” (black solid line) spike during high pollution events (dashed line), particularly the major wildfire episodes of August-September 2020. Panel B demonstrates that people not only search for “purple air” (line), but also follow through, as the searches strongly predict new monitor installations (bars) in our data, with a correlation of 0.73 at weekly frequencies (see Coury et al. (2024) for causal evidence on monitor adoption in response to pollution spikes). While we do not directly observe air purifier purchases, this pattern is strongly suggestive that households respond to pollution episodes by acquiring both monitoring and protective technologies.

2.3 Data and Sample Construction

Our analysis uses data from all indoor PurpleAir monitors installed in the U.S. between January 2019 and October 15, 2021.³ We identify indoor monitors using the “location_type” field in the PurpleAir database. For each monitor, we downloaded the hourly average $\text{PM}_{2.5}$ concentrations for analysis. To ensure data quality, we follow the manufacturer’s recommended filtering procedures.⁴

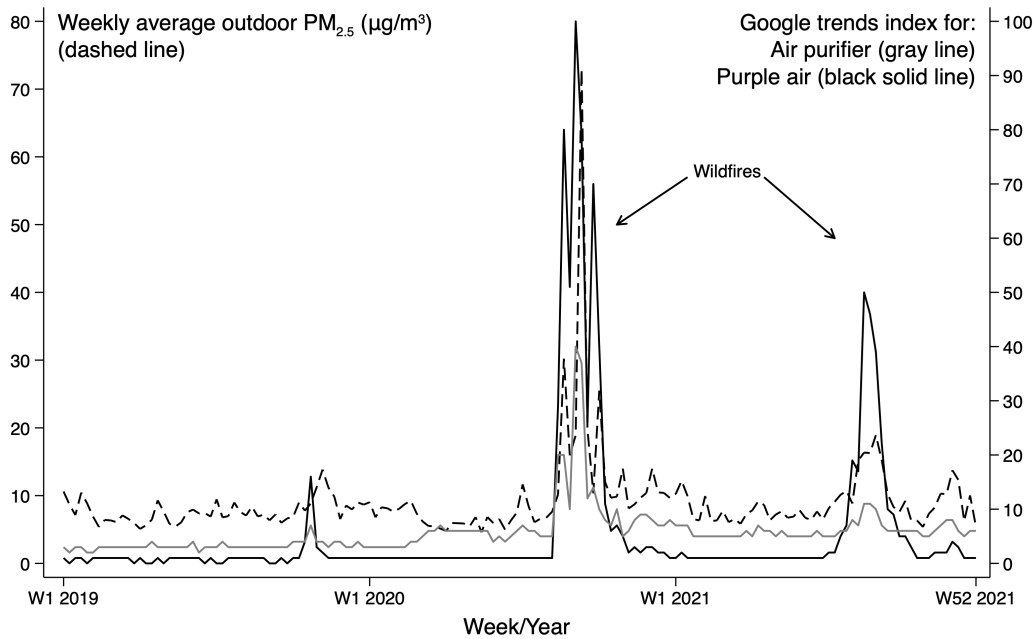
We restrict the sample to monitors with clearly identified installation times and sufficiently complete post-installation observations according to the following steps. First, we drop monitors that disappear within the first twelve weeks (no observations beyond hour 2016). Second, to identify the proper date and hour of monitor installation, we distinguish between two installer types based on broadcasting patterns. Type 1 monitors broadcast uninterruptedly during all hours of the first week after going online. Our hourly count for type 1 monitors starts with the first observation. Type 2

³We include only monitors installed by October 15, 2021, to ensure twelve weeks of observation before the end of the sample period, which is December 31, 2021.

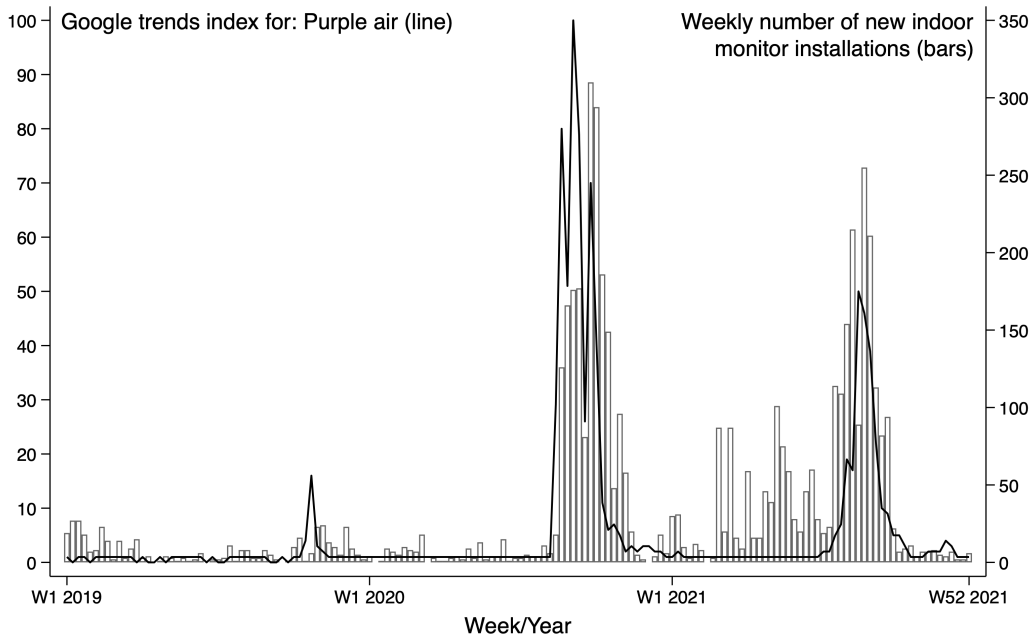
⁴See <http://www.aqmd.gov/docs/default-source/aq-spec/resources-page/>, accessed May 5, 2020. As recommended, we exclude observations where the two measurement channels (A and B) differ by more than $10 \mu\text{g}/\text{m}^3$ for concentrations below $100 \mu\text{g}/\text{m}^3$ or by more than 10% for concentrations above that threshold. This filter can be applied only to the subset of devices with dual measurements. In addition, we drop all observations exceeding $500 \mu\text{g}/\text{m}^3$, which likely reflect sensor malfunction or contamination.

Figure 1: Ambient pollution correlation with monitoring and defense interest

Panel A: Outdoor air pollution vs. Google searches



Panel B: “Purple air” Google searches vs. PurpleAir indoor monitor installations



Panel A: Weekly average outdoor PM_{2.5} (µg/m³) (dashed line, left y-axis) and weekly Google Trends Index for “air purifier” and “purple air” (gray line and black solid line, right y-axis) by weeks of 2019-2021 (x-axis) in California. **Panel B:** Weekly Google Trends Index for “purple air” (line, left y-axis) and weekly number of new PurpleAir indoor monitor installations (bars, right y-axis) by weeks of 2019-2021 (x-axis) in California. **Notes:** The Google Trends Index is scaled relative to peak search interest (100 for “purple air” in September 2-8, 2020).

monitors experience an interruption of at least 24 hours within the first week after first going online, but then broadcast uninterrupted during all hours of the first week following that interruption. Our hourly count for type 2 monitors starts with the first observation after the interruption, and we drop all observations before the interruption. We further drop all monitors that do not match one of these types. This step is necessary because many monitors briefly go online, presumably upon delivery, when users test and connect the device through the app, before being temporarily disconnected until they find time to properly wall-mount the monitor. Our approach ensures that the identified installation time reflects the moment households begin continuously receiving IAQ information, rather than an initial setup test. 87% of our final sample are type 2 monitors.⁵ Third, we only keep monitors without any missing observations within the first 24 hours and at least 90% non-missing observations in the twelve weeks after installation. This includes missing outdoor pollution and weather data merged based on the Census tract, as described below. Finally, we drop all observations beyond week 12. Our final sample includes 4,470,924 hourly observations from 2,249 monitors in 1,340 Census tracts. 83% of monitors are in California. The next two states are Washington (4%) and Oregon (3%).

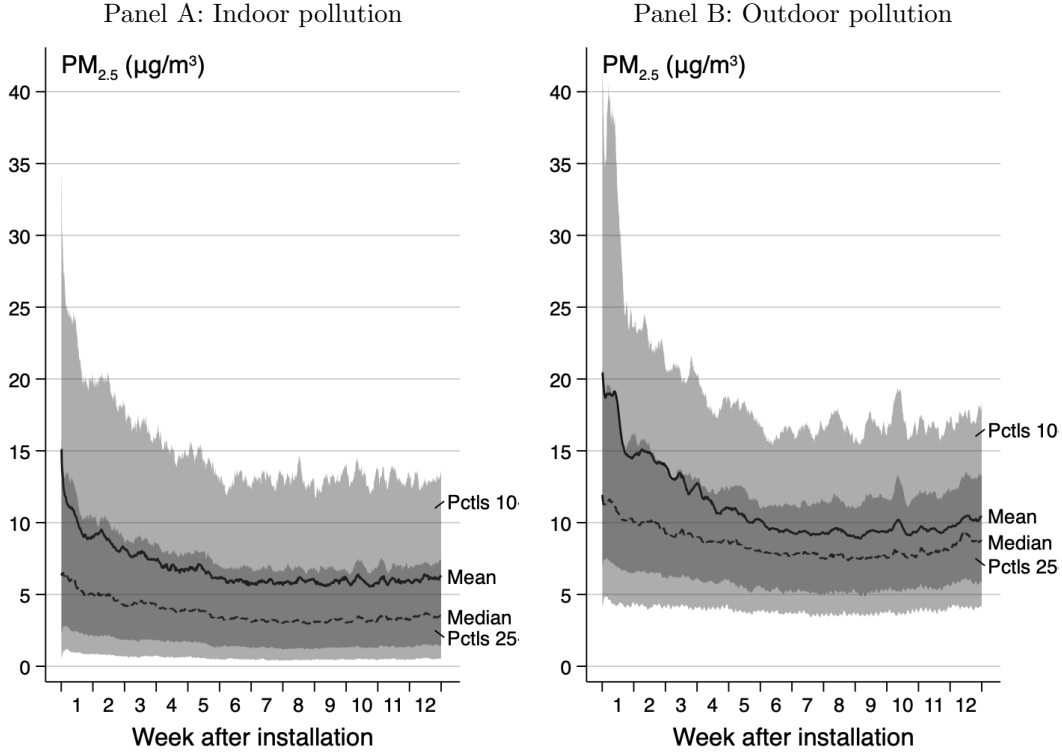
We match each indoor monitor to outdoor pollution data from the EPA Air Quality System and to weather information from the NOAA Global Hourly Integrated Surface Database based on the indoor monitor’s Census tract, inferred from its reported coordinates.⁶ The weather variables include air temperature, dew point temperature, precipitation, and wind speed.

Figure 2 presents the raw evolution of indoor pollution (Panel A) and the corresponding outdoor pollution (Panel B) across all indoor monitors after their installation, with the solid line representing the mean, the dashed line the median, the darker shaded area the interquartile range, and the lighter shaded area the range from the 10th to the 90th percentiles of PM_{2.5} concentrations. Panel A shows that the mean indoor PM_{2.5} starts at approximately 15 $\mu\text{g}/\text{m}^3$ on the installation day and declines to 6 $\mu\text{g}/\text{m}^3$ after five weeks and stays at the same level through week 12. Panel B reveals that outdoor pollution follows a similar pattern, falling from 20 to 9 $\mu\text{g}/\text{m}^3$ by week 7, with a slight uptick towards the end of the same period. This concurrent movement, due to more monitor purchases when outdoor pollution is high, as shown in Figure 1, underscores the importance of controlling for outdoor conditions.

⁵Because type 2 monitors may receive information prior to the installation time we define, our estimates may represent a lower bound of the true effect. We also replicate all analyses using only type 1 monitors. While confidence intervals are wider and some estimates lose statistical significance, the qualitative pattern of results is unchanged.

⁶For both EPA and NOAA data, we match monitors to Census tract centroids using inverse-distance-squared weighting of all monitors within 50 miles. We retain only monitors with at least 2/3 non-missing observations. For weather data, we impute missing values using the inverse-distance-squared weighted average of other monitors within 50 miles before constructing tract-level averages. EPA AQS data: https://aqs.epa.gov/aqsweb/airdata/download_files.html#Raw, accessed August 21, 2024. NOAA ISD data: <https://www.ncei.noaa.gov/products/land-based-station/integrated-surface-database>, accessed August 21, 2024.

Figure 2: Indoor and outdoor air pollution reduction over time



Panel A: Mean (solid line), median (dashed line), percentiles 25-75 (darker shaded area), and percentiles 10-90 (lighter shaded area) of indoor $\text{PM}_{2.5}$ ($\mu\text{g}/\text{m}^3$) during the first twelve weeks after indoor monitor installation, as measured by all indoor monitors. **Panel B:** Mean (solid line), median (dashed line), percentiles 25-75 (darker shaded area), and percentiles 10-90 (lighter shaded area) of outdoor $\text{PM}_{2.5}$ ($\mu\text{g}/\text{m}^3$) during the first twelve weeks after indoor monitor installation, as measured by the EPA monitoring network in the corresponding Census tracts. **Notes:** We first calculate 24-hour moving averages to then compute means, medians, and percentiles.

2.4 Identifying Household Behaviors

A key challenge in our setting is that we observe pollution outcomes but not the behaviors that generate them. We address this by developing machine learning algorithms to identify two important household responses: cooking activity and air purifier adoption.

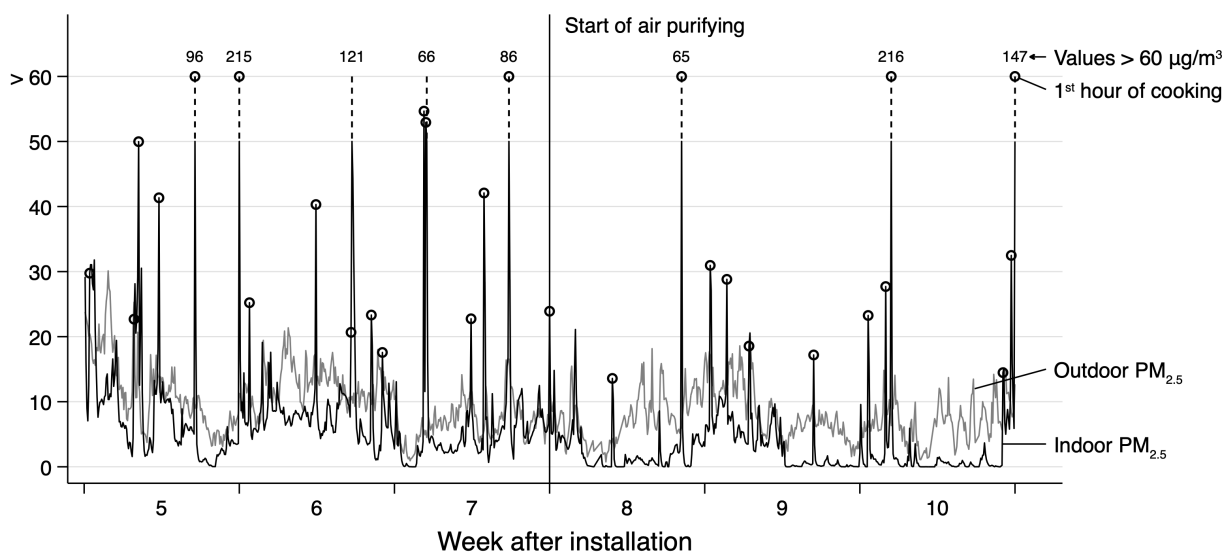
2.4.1 Cooking Episodes

Cooking generates distinctive spikes in indoor $\text{PM}_{2.5}$ through particles released when oils and foods are heated (Torkmahalleh et al. 2017; Lachowicz et al. 2022). We identify cooking episodes based on four steps. First, we calculate the residuals from separate linear regressions for each indoor monitor and week since installation of indoor on outdoor $\text{PM}_{2.5}$. This ensures we do not classify indoor spikes as induced by outdoor pollution. We separate by monitor and week since indoor infiltration of outdoor pollution varies both by household and over time, as people adopt mitigation strategies. Second, we run the Isolation Forest algorithm (Liu et al. 2008) as implemented in the scikit-learn Python package (Pedregosa et al. 2011) to detect anomalies in the hour-by-hour changes of these

residuals. Third, we condition outliers on being positive and higher than $10 \mu\text{g}/\text{m}^3$, and classify them as the start of a cooking episode. Fourth, we count the six hours from the initial spike as the cooking episode. While typical cooking activities last less than six hours, our cooking episode aims to capture the period of elevated pollution following the initial spike. As shown below, pollution from cooking often persists beyond six hours, meaning our definition is conservative.

Figure 3 illustrates this approach for a representative household. The black line represents the hourly indoor $\text{PM}_{2.5}$ readings for weeks five to ten after the monitor installation. Values above $60 \mu\text{g}/\text{m}^3$ are labeled with the peak value (and the black line becomes dashed above $50 \mu\text{g}/\text{m}^3$). The gray line shows the corresponding outdoor $\text{PM}_{2.5}$ concentration. Circles indicate initial cooking spikes or the start of a cooking episode. The figure reveals many sharp spikes in indoor $\text{PM}_{2.5}$ that do not correspond to outdoor changes, supporting our detection approach.

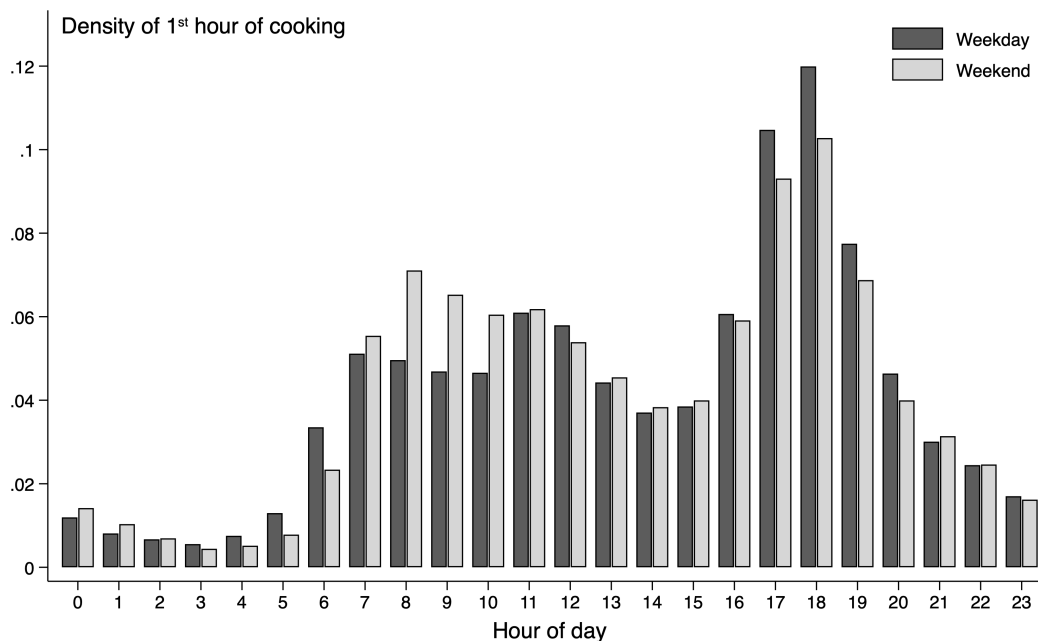
Figure 3: Example household



Hourly average readings of indoor $\text{PM}_{2.5}$ ($\mu\text{g}/\text{m}^3$) (black line) and outdoor $\text{PM}_{2.5}$ ($\mu\text{g}/\text{m}^3$) in the corresponding Census tract (gray line) of an example household during weeks 5 to 10 after indoor monitor installation, as well as hours coded as 1st hour of a cooking episode (circles). The vertical line marks the start of air purifying for this example household, as detected by our algorithm. **Notes:** Values above $60 \mu\text{g}/\text{m}^3$ are shown as 60 and labeled with their corresponding actual values, rounded to the nearest integer. These peaks mostly coincide with the 1st hour of a cooking episode.

Figure 4 further validates our cooking detection algorithm. It shows the timing (density) of identified initial cooking spikes, separately for weekdays and weekends. The start of cooking episodes starkly rises in the morning and again in the early evening hours. Cooking starts later, and happens more in morning hours and less in evening hours on weekends, matching expected meal preparation patterns. Note that our sample period includes the COVID-19 pandemic, when cooking frequency and timing likely differed from normal due to increased time at home.

Figure 4: Cooking hours



Density of 1st hour of a cooking episode by hour of day for weekdays (darker bars) and weekend days (lighter bars), as detected by our algorithm.

2.4.2 Air Purifier Adoption

Air purifiers reduce indoor $\text{PM}_{2.5}$ by filtering particles from recirculated air. To identify purifier usage, we focus on indoor pollution levels, excluding cooking episodes. This exclusion is essential because cooking generates temporary pollution spikes regardless of purifier operation (see, e.g., Xiang et al. (2021)), which would obscure the persistent changes in baseline pollution that indicate continuous filtration. For each household-week, we calculate the share of non-cooking hours where indoor $\text{PM}_{2.5}$ remains below $3 \mu\text{g}/\text{m}^3$ among hours with outdoor $\text{PM}_{2.5}$ above $3 \mu\text{g}/\text{m}^3$ —a threshold indicating near-complete particle removal achieved by effective purifier operation.⁷ Finally, we apply the k-Means clustering algorithm (MacQueen 1967; Lloyd 1982) as implemented in the scikit-learn library (Pedregosa et al. 2011) for cluster analysis of this share.

We aggregate to the weekly level rather than the daily level because indoor pollution can be temporarily elevated on any given day for reasons unrelated to purifier usage, such as extensive cooking, open windows during high outdoor pollution, or other transient household activities.

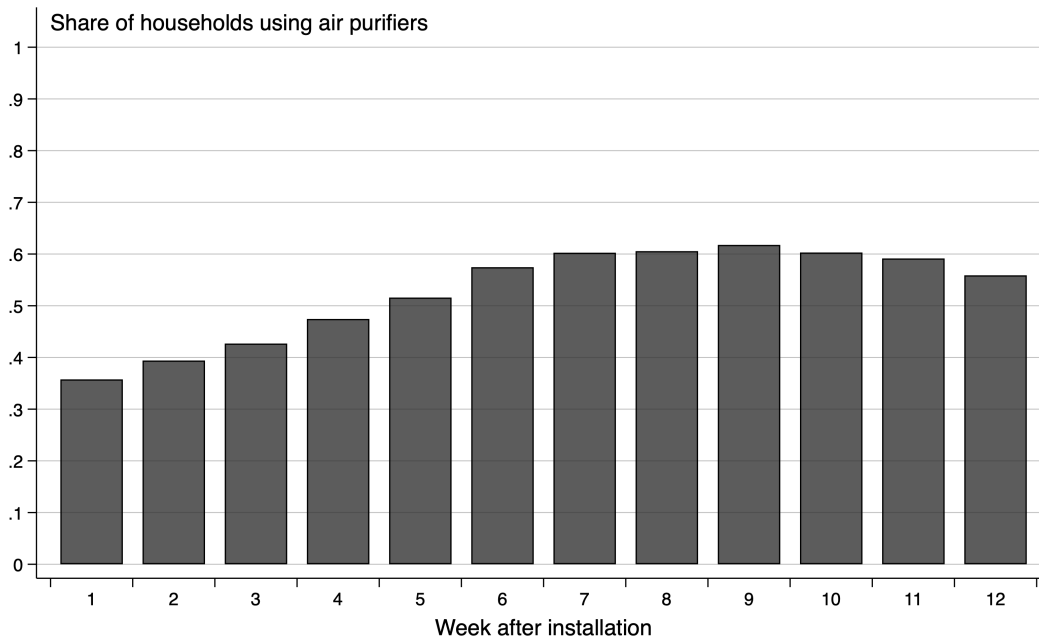
⁷HEPA filters are defined in U.S. regulation as removing at least 99.97% of particles $\geq 0.3 \mu\text{m}$ in a single pass (<https://www.epa.gov/indoor-air-quality-iaq/what-hepa-filter?>, accessed November 19, 2025). Under continuous recirculation, repeated passes through high-efficiency filters produce order-of-magnitude reductions in particle concentrations. Residential intervention studies from the U.S. and Canada confirm this (for recent literature reviews, see Allen and Barn (2020), Cheek et al. (2021), Zhu et al. (2021)). Most studies, while not excluding cooking hours, find post-purifying average $\text{PM}_{2.5}$ values below $5 \mu\text{g}/\text{m}^3$, with the lowest reporting a mean of $4 \mu\text{g}/\text{m}^3$ and a median of $2.6 \mu\text{g}/\text{m}^3$ for residences in Fresno, California (Cheng et al. 2016). We therefore chose $3 \mu\text{g}/\text{m}^3$ as a conservative approximation to purified background pollution during non-cooking hours. Indoor values below $3 \mu\text{g}/\text{m}^3$ when outdoor $\text{PM}_{2.5}$ is above $3 \mu\text{g}/\text{m}^3$ are unlikely to occur without active filtration.

Daily or sub-daily classification would misattribute these short-term fluctuations to changes in purifier behavior. Weekly aggregation captures sustained patterns in baseline pollution that reliably indicate whether a household is consistently operating filtration equipment. Based on this weekly measure, we observe little on-off switching: among households that adopt purifiers, the typical household switches its purifier off once and back on once during our twelve-week observation period.

Returning to [Figure 3](#), we can see the effect of purifier adoption for the example household. Before the 8th week after monitor installation, indoor and outdoor pollution are highly correlated. During week 8, the relationship between indoor and outdoor pollution changes, with indoor levels consistently below outdoor levels and in many hours close to 0. Although this household may not have used its purifier consistently after adoption, our algorithm classifies all weeks from week 8 as purifier usage.

[Figure 5](#) shows that purifier adoption increases gradually over time, from 36% to 62% of households by week 9. The gradual adoption pattern is consistent with households needing time to research options, make purchasing decisions, and receive deliveries, rather than responding immediately to their first pollution reading. The fact that it stays constant after 9 weeks suggests that additional time and information appear unnecessary for deciding to invest in a purifier.

Figure 5: Air purifier adoption



Share of households using air purifiers by week after indoor monitor installation, as detected by our algorithm.

3 Empirical Strategy

3.1 Main Specification

Our baseline specification examines changes in indoor PM_{2.5} following monitor installation while flexibly controlling for contemporaneous outdoor pollution and weather:

$$I_{mh} = \sum_{d=2}^7 \beta_d^1 1(D_{mh} = d) + \sum_{w=2}^{12} \gamma_w^1 1(W_{mh} = w) + \mathbf{O}'_{mh} \delta^1 + \mathbf{X}'_{mh} \eta^1 + \iota_m^1 + \lambda_h^1 + \varepsilon_{mh}^1, \quad (1)$$

where I_{mh} is the indoor PM_{2.5} concentration for monitor m in hour h , and D_{mh} and W_{mh} count days 2-7 and weeks 2-12 since installation. In terms of control variables, \mathbf{O}_{mh} is a vector of outdoor PM_{2.5} pollution bin indicators⁸, \mathbf{X}_{mh} includes weather controls⁹, ι_m^1 denotes monitor fixed effects, and λ_h^1 captures state-year-month, day-of-week, and hour-of-day fixed effects. Notably, the monitor fixed effects ensure we are comparing changes in IAQ over time within the same household, thereby controlling for all time-invariant characteristics of the household. ε_{mh}^1 is the error term. The coefficients of interest, β_d^1 and γ_w^1 , capture changes in IAQ for days 2-7 and weeks 2-12 relative to the first 24 hours after installation. Coefficient superscripts denote the equation number in all equations.

The key identifying assumption for interpreting our estimates as causal effects of information on indoor pollution requires that households' first monitor reading represents genuinely new information. That is, households had limited or inaccurate priors about IAQ before installation. This assumption seems reasonable—people purchase monitors specifically to learn about their IAQ—but it may not hold perfectly. For example, households may have had prior expectations about IAQ based on observable cues, such as odors or visible smoke, and already knew their air quality was problematic. While we cannot verify this directly, we suspect this is unlikely because PM_{2.5}, our measure of IAQ, is correlated with other observable cues but is itself undetectable by human senses. This suggests that any knowledge of baseline pollution levels is likely quite noisy. To the extent some households did know and were already mitigating poor IAQ, our estimates would understate the effect of information.

The observed temporal evolution of IAQ following monitor installation (Equation 1) helps to assess this concern. If households possessed substantial prior knowledge about their IAQ and pre-determined mitigation plans, we would expect immediate improvements following installation, as households quickly implement their planned strategies. In contrast, a true information-response mechanism should produce gradual improvements as households require time to understand, process, and translate information from monitor readings into effective mitigation strategies.

While there may be a delay from households acquiring the information to their reaction, some

⁸The bins for outdoor PM_{2.5} concentration are steps of 5 from 0 to 100 $\mu\text{g}/\text{m}^3$, steps of 10 from 100 to 200 $\mu\text{g}/\text{m}^3$, steps of 50 from 200 to 300 $\mu\text{g}/\text{m}^3$, and the highest bin indicating readings higher than 300 $\mu\text{g}/\text{m}^3$.

⁹Air and dew point temperatures, air and dew point temperatures squared (with separate variables for positive and negative temperatures), and second-order polynomials for wind speed and precipitation.

households may already engage in mitigation within our reference periods: 24 hours for our main results (Section 4.1 & Section 4.2) and one week for our indoor-outdoor infiltration (Section 4.3) and cooking frequency (Section 5.1). To the extent adaptation occurs as a result of the newly provided information within these reference periods, our estimates will underestimate the total impact.

To further strengthen causal claims, we also explore heterogeneity in responses based on initial indoor and outdoor pollution levels. If households are responding to new information rather than acting on prior knowledge, we should observe differential responses based on the content of that information. Households discovering their IAQ is relatively clean have limited scope and incentive for improvement. In contrast, households discovering problematic IAQ have both greater room for improvement and higher marginal health benefits from mitigation. We explore these impacts by stratifying Equation 1 by initial indoor and outdoor levels.

Another concern with interpreting our results as causal is that monitor adoption coincides with high pollution episodes. Purchases spike during wildfires and other pollution events, creating a mechanical tendency for air quality to improve from these abnormal starting points, even absent any behavioral response. Given the high indoor infiltration rates of PM, declining outdoor conditions following peak episodes would mechanically reduce indoor levels.

To address these concerns, we control flexibly for contemporaneous outdoor air quality through \mathbf{O}_{mh} in all specifications (see Footnote 8). This approach accounts for the mechanical relationship between outdoor and indoor pollution, ensuring that improvements in IAQ are not simply reflecting declining outdoor conditions. However, we recognize this may not fully capture all mechanical trends, particularly if indoor pollution responds to outdoor conditions with lags or if \mathbf{O}_{mh} does not fully capture the variation in the indoor-outdoor relationship during extreme episodes. Therefore, we further probe this by investigating changes over time in the indoor-outdoor gradient.

3.2 Indoor-Outdoor Gradient

Our main specification (Equation 1) examines how indoor pollution changes over time conditional on outdoor pollution levels. An equally important question is whether households become better at protecting themselves, specifically during high outdoor pollution episodes, when health damages are most severe. If information primarily helps households avoid exposure during these critical periods, the health benefits from monitor adoption would be substantially larger than average reductions suggest.

To explore this, we examine how the relationship between indoor and outdoor pollution evolves following monitor installation. A declining gradient over time, particularly at high outdoor pollution levels, indicates that households are increasingly able to maintain clean indoor air despite poor outdoor conditions, consistent with defensive investments like air purifiers.

Examining indoor pollution reductions across outdoor pollution levels provides additional evidence against mechanical mean reversion due to declining outdoor levels. Rather than simply controlling for outdoor conditions, we test whether households achieve lower indoor pollution at specific outdoor levels over time, addressing concerns that flexible controls may inadequately cap-

ture outdoor pollution’s mechanical effects.

Our estimation model interacts time since monitor installation with outdoor pollution bin indicators to quantify any differences in the indoor-outdoor gradient over time:

$$I_{mh} = \sum_{s=2}^4 \alpha_s^2 1(S_{mh} = s) + \sum_{b=2}^8 \beta_b^2 1(B_{mh} = b) + \sum_{s=2}^4 \sum_{b=2}^8 \gamma_{sb}^2 1(S_{mh} = s) 1(B_{mh} = b) + \mathbf{X}'_{mh} \eta^2 + \iota_m^2 + \lambda_h^2 + \varepsilon_{mh}^2. \quad (2)$$

The time variable, S_{mh} , distinguishes three post-monitor installation periods ($S_{mh} = 2$: weeks 2-4; $S_{mh} = 3$: weeks 5-8; $S_{mh} = 4$: weeks 9-12), with week 1 as the reference. The outdoor pollution bin variable, B_{mh} , represents seven bins in $5 \mu\text{g}/\text{m}^3$ increments ($B_{mh} = 2$: outdoor $\text{PM}_{2.5} = 5\text{-}10$, continuing through $B_{mh} = 7$: outdoor $\text{PM}_{2.5} = 30\text{-}35 \mu\text{g}/\text{m}^3$, and $B_{mh} = 8$: outdoor $\text{PM}_{2.5} \geq 35 \mu\text{g}/\text{m}^3$), with $\text{PM}_{2.5} \leq 5 \mu\text{g}/\text{m}^3$ as the reference bin. All controls are defined analogously to Equation 1. The coefficients α_s^2 capture indoor pollution changes over time at the reference outdoor pollution level, and the interaction coefficients, γ_{sb}^2 , measure how the indoor-outdoor gradient changes over time, testing whether households become increasingly effective at lowering IAQ as outdoor pollution rises.

3.3 Cooking Behavior

To investigate changes in cooking behavior, we aggregate our model to the week level and estimate the following Poisson pseudo-maximum likelihood model for weekly cooking events:

$$C_{mw} = \exp \left(\sum_{w=2}^{12} \gamma_w^3 1(W_{mw} = w) + \mathbf{O}'_{mw} \delta^3 + \mathbf{X}'_{mw} \eta^3 + \iota_m^3 + \lambda_w^3 \right). \quad (3)$$

Here, C_{mw} is the number of cooking events detected by monitor m in week w . \mathbf{O}'_{mw} and \mathbf{X}'_{mw} are weekly averages of outdoor $\text{PM}_{2.5}$ and variables constructed from weekly averages of the same weather variables used in all regressions. Given the aggregation to the weekly level, λ_w^3 only includes state-year-month fixed effects but not day-of-week and hour-of-day fixed effects. The main coefficient of interest, γ_w^3 , indicates the change in cooking frequency in week w , relative to week 1.

We conduct the analysis at the weekly level because our outcome variable—the number of cooking events—is a count variable. The main limitation of this approach is that some behavioral adjustment may already occur during the first week after monitor installation, which serves as the reference period. Consequently, the estimated effects are to be interpreted as a lower bound.

3.4 Ventilation Behavior

To explore whether people ventilate more when cooking, we restrict our sample to households that never adopt an air purifier or household-weeks occurring before purifier adoption. Further, we only

keep households with at least twelve cooking events before purifier adoption (or over the full sample period for non-adopters). We then estimate the following model:

$$I_{mh} = \sum_{n=2}^4 \alpha_n^4 1(N_{mh} = n) + \sum_{c=1}^6 \beta_c^4 1(C_{mh} = c) + \sum_{n=2}^4 \sum_{c=1}^6 \gamma_{nc}^4 1(N_{mh} = n) 1(C_{mh} = c) + \mathbf{O}'_{mh} \delta^4 + \mathbf{X}'_{mh} \eta^4 + \iota_m^4 + \lambda_h^4 + \varepsilon_{mh}^4. \quad (4)$$

N_{mh} captures bins for the n -th cooking event after monitor installation: $N_{mh} = 2$: 2nd or 3rd event; $N_{mh} = 3$: 4th to 6th event; $N_{mh} = 4$: 7th to 12th event. We drop observations after the 12th event. $C_{mh} = \{1, \dots, 6\}$ counts the hours since the start of the cooking event. Non-cooking hours before the 2nd cooking event serve as the reference category. All controls are defined analogously to Equation 1. Thus, α_n^4 captures the change in indoor pollution during non-cooking times across cooking-event bins n , β_c^4 the temporal pattern of the first cooking spike across hours c , and the interaction terms, γ_{nc}^4 , how pollution spikes evolve across cooking-event bins n and hours c .

3.5 Air Purifier Effectiveness

To explore air purifier effectiveness in reducing the indoor-outdoor gradient, we estimate the following model:

$$I_{mh} = \alpha^5 F_{mh} + \sum_{b=2}^8 \beta_b^5 1(B_{mh} = b) + \sum_{b=2}^8 \gamma_b^5 F_{mh} 1(B_{mh} = b) + \mathbf{X}'_{mh} \eta^5 + \iota_m^5 + \lambda_h^5 + \varepsilon_{mh}^5. \quad (5)$$

F_{mh} indicates filtration adoption. B_{mh} represents the same seven outdoor pollution bins as in Equation 2. All controls are defined analogously to Equation 1. β_b^5 is the indoor-outdoor gradient for those who do not adopt purifiers, and γ_b^5 captures the change in the gradient from air purifying.

Similarly, to analyze air purifier effectiveness during cooking, we estimate the following model, again only including households with at least twelve cooking events:

$$I_{mh} = \alpha^6 F_{mh} + \sum_{c=1}^6 \beta_c^6 1(C_{mh} = c) + \sum_{c=1}^6 \gamma_c^6 F_{mh} 1(C_{mh} = c) + \mathbf{O}'_{mh} \delta^6 + \mathbf{X}'_{mh} \eta^6 + \iota_m^6 + \lambda_h^6 + \varepsilon_{mh}^6 \quad (6)$$


Again, F_{mh} indicates filtration adoption and C_{mh} , as in Equation 4, counts the hours since the start of the cooking event. All controls are defined analogously to Equation 1. α^6 captures the change in indoor pollution during non-cooking hours and γ_c^6 how pollution spikes change with purifying across hours c since the start of the cooking episode.

4 Changes in Indoor Air Quality

4.1 Main Results

Figure 6 presents our main results on changes in indoor $\text{PM}_{2.5}$ following monitor installation. The figure plots coefficients and 95% confidence intervals from Equation 1. Indoor pollution decreases steadily over time, with reductions becoming statistically significant by the second day and growing to approximately $2.5 \mu\text{g}/\text{m}^3$ by week 12. Interpreting the magnitude of this impact requires nuance because the appropriate baseline is not obvious. Comparing to the first-day mean (11.76) yields a 21.3% decrease, but this likely underestimates the true effect since IAQ is typically lower than first-day readings. Conversely, comparing to the overall mean (6.76) yields a 37% decrease, but this overestimates the effect by double-counting the monitor-induced reduction already present in the baseline. The total reduction from a raw comparison between day 1 to week 12 is $5.6 \mu\text{g}/\text{m}^3$, suggesting mitigation accounted for about 45% of the total reduction.

Figure 6: Estimated change in indoor pollution by time since monitor installation



Coefficients with 95% confidence intervals from day and week after indoor monitor installation indicators (β_d^1 and γ_w^1 in Equation 1). **Notes:** The regression includes flexible controls for outdoor pollution and weather as well as indoor monitor, state-year-month, day-of-week, and hour-of-day fixed effects. Standard errors are clustered on Census tracts.

The gradual nature of improvements over twelve weeks supports an information-based mechanism. If households had prior knowledge and were simply waiting for monitors to confirm it, we might expect more immediate responses. Instead, we observe gradual, continuous improvements over the full twelve-week period, with effects growing steadily from week 1 through week 12, suggesting a learning and adjustment process. Furthermore, if improvements were purely mechanical artifacts of monitors being purchased during peak pollution episodes, we would expect the steepest

improvements in the first few weeks as conditions naturally normalize, followed by a plateau. The sustained, linear improvement trajectory we document is difficult to reconcile with such mechanical explanations but aligns well with households progressively responding to new information.

4.2 Heterogeneous Effects by Initial Conditions

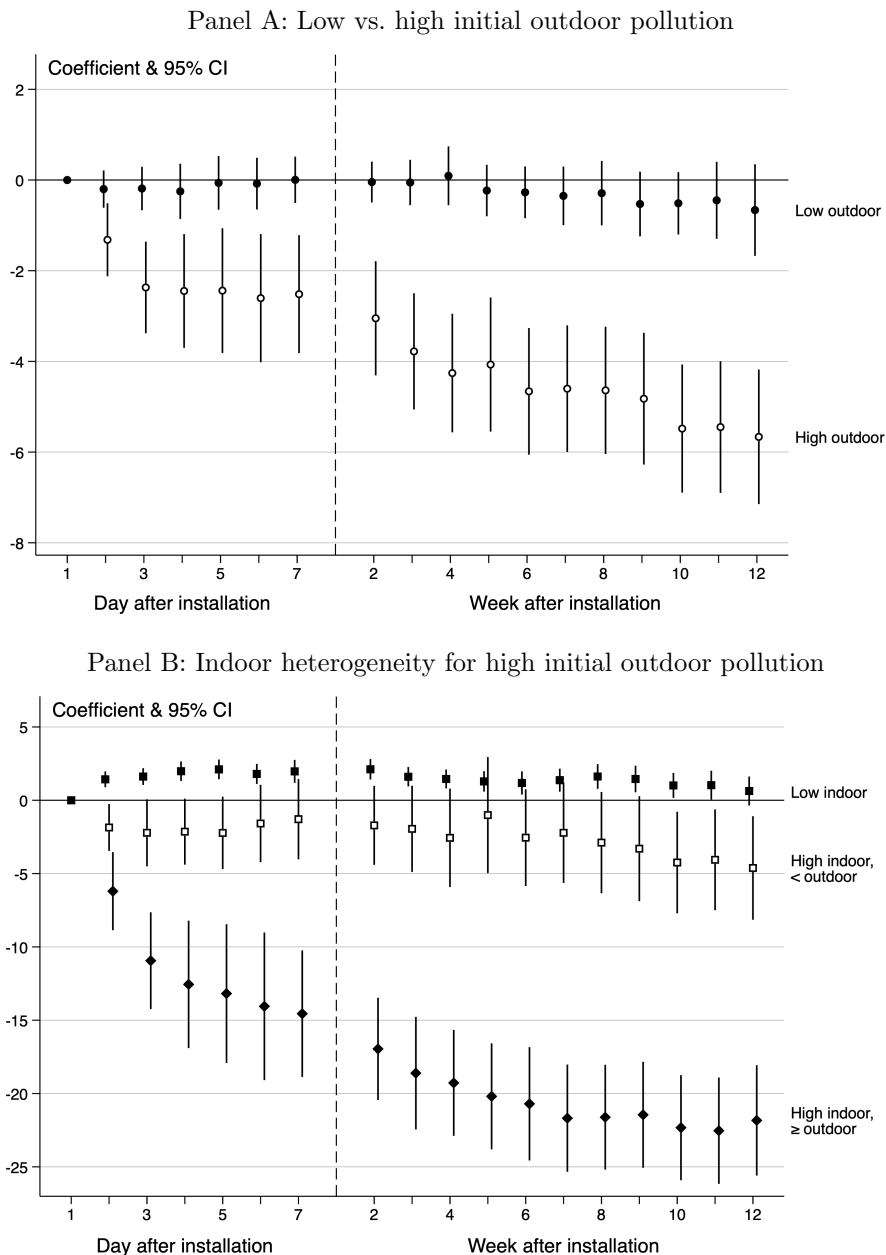
To further probe these results, we examine heterogeneous effects based on initial pollution levels. We first stratify our sample based on whether first-day mean outdoor $\text{PM}_{2.5}$ is below or above $12 \mu\text{g}/\text{m}^3$ (as before, we call this low vs. high). Panel A of Figure 7 shows the overall effect for these two subsamples, with essentially no change in IAQ for households with low initial outdoor concentrations over twelve weeks (filled circles). Given that these households' mean first-day indoor $\text{PM}_{2.5}$ level is $5.4 \mu\text{g}/\text{m}^3$, they have very limited scope for gains. In contrast, households with high initial outdoor pollution (hollow circles)—starting from a first-day mean indoor $\text{PM}_{2.5}$ level of $18.9 \mu\text{g}/\text{m}^3$ —experience considerable improvements over time by almost $6 \mu\text{g}/\text{m}^3$ by week 12, with a sizable decrease already occurring on day 2.

Panel B of Figure 7 further examines households with high initial outdoor $\text{PM}_{2.5}$, splitting the sample by initial indoor levels, now with three categories: low first-day indoor $\text{PM}_{2.5}$, high first-day $\text{PM}_{2.5}$ but lower than outdoor $\text{PM}_{2.5}$, and high first-day indoor $\text{PM}_{2.5}$ that is even higher than (high) outdoor $\text{PM}_{2.5}$. Exploring separately by indoor relative to outdoor $\text{PM}_{2.5}$ allows us to assess the role of ventilation available due to better outdoor air quality relative to IAQ.

Households with low initial indoor pollution (on average $6.4 \mu\text{g}/\text{m}^3$) show no improvement (filled squares), sensibly, since increased ventilation would worsen IAQ. In fact, we observe a slight increase in indoor pollution during the first week, which reverts to initial levels over time. Households with high initial indoor pollution (on average $25.8 \mu\text{g}/\text{m}^3$) that remains below outdoor levels show modest improvements beginning in week 2 (hollow squares). Since ventilation would worsen their IAQ, these households cannot respond immediately. The delayed, gradual improvements are suggestive of defense investments that take time to acquire, such as air purifiers. Most dramatically, households where initial indoor pollution is high (on average $39.8 \mu\text{g}/\text{m}^3$) and exceeds outdoor pollution show both an immediate drop during week 1 and sustained improvements thereafter, reaching a reduction of $22 \mu\text{g}/\text{m}^3$ by week 12 (diamonds). These households can immediately improve their IAQ through ventilation, reducing indoor pollution down to outdoor levels. The continued decline beyond week 1 suggests subsequent adoption of defense investments like air purifiers.

These differential patterns—improvements concentrated where indoor exceeds outdoor pollution, regardless of absolute levels—suggest purposeful household mitigation rather than mechanical mean reversion. Moreover, the evidence that households only improve IAQ when it makes sense given their specific circumstances (households only respond when indoor levels are high) and when it is practicable (immediate reductions, presumably from opening windows, only occur in households where indoor levels exceed outdoor levels) is more consistent with households learning about and responding to their specific situation than with pre-existing plans.

Figure 7: Estimated change in indoor pollution by time since monitor installation, heterogeneity by initial outdoor and initial indoor pollution




Panel A: Coefficients with 95% confidence intervals from day and week after indoor monitor installation indicators (β_d^1 and γ_w^1 in Equation 1), from separate regressions for indoor monitors with average first-day outdoor $\text{PM}_{2.5} < 12 \mu\text{g}/\text{m}^3$ (filled circles) and $\geq 12 \mu\text{g}/\text{m}^3$ (hollow circles). **Panel B:** Coefficients with 95% confidence intervals from day and week after indoor monitor installation indicators (β_d^1 and γ_w^1 in Equation 1), from separate regressions for indoor monitors with average first-day indoor $\text{PM}_{2.5} < 12 \mu\text{g}/\text{m}^3$ (filled squares) vs. average first-day indoor $\text{PM}_{2.5} \geq 12 \mu\text{g}/\text{m}^3$ & $<$ first-day outdoor $\text{PM}_{2.5}$ (hollow squares) vs. average first-day indoor $\text{PM}_{2.5} \geq 12 \mu\text{g}/\text{m}^3$ & \geq first-day outdoor $\text{PM}_{2.5}$ (diamonds), restricted to the sample of indoor monitors with average first-day outdoor $\text{PM}_{2.5} \geq 12 \mu\text{g}/\text{m}^3$. **Notes:** All regressions include flexible controls for outdoor pollution and weather as well as indoor monitor, state-year-month, day-of-week, and hour-of-day fixed effects. Standard errors are clustered on Census tracts.

4.3 Evolution of Indoor-Outdoor Gradient

As previously mentioned, we explore how the indoor-outdoor gradient evolves to understand how people learn over time, as IAQ may change. Figure 8 provides a first glimpse into this relationship by plotting the raw data. The figure depicts mean indoor (y-axis) and outdoor pollution (x-axis) across 100 percentiles of outdoor pollution for weeks 1 (filled circles), 2-4 (hollow circles), 5-8 (filled squares), and 9-12 (hollow squares) after installation, along with fitted values from separate 3rd-order polynomial regression of indoor on outdoor pollution for the same time periods (direct labeling in the figure).¹⁰ The gradient steadily flattens, with small differences at low outdoor pollution but large improvements at high levels.

Figure 8: Changing indoor-outdoor pollution relationship



Mean indoor and outdoor $\text{PM}_{2.5}$ in $\mu\text{g}/\text{m}^3$ for 100 percentiles of the distribution of outdoor $\text{PM}_{2.5}$ values in week 1 (filled circles), weeks 2-4 (hollow circles), weeks 5-8 (filled squares), and weeks 9-12 (hollow squares) after installation. Fitted values of a 3rd-order polynomial regression of indoor on outdoor pollution, separately for week 1, weeks 2-4, weeks 5-8, and weeks 9-12 after installation. The corresponding lines are labeled directly in the figure. **Notes:** The data for the regression is trimmed at the outdoor $\text{PM}_{2.5}$ level of $34.46 \mu\text{g}/\text{m}^3$, equivalent to the 99th percentile of the distribution of outdoor $\text{PM}_{2.5}$ values in weeks 9-12. Accordingly, the figure is trimmed at the outdoor $\text{PM}_{2.5}$ level of $34.46 \mu\text{g}/\text{m}^3$.

Figure 9 provides results from estimating Equation 2, plotting predicted indoor $\text{PM}_{2.5}$ based on joint estimates of the relevant parameters. Each point in the figure reflects a linear combination of coefficients. For a given post-monitor installation period, S_{mh} , and outdoor pollution bin, B_{mh} , predicted indoor $\text{PM}_{2.5}$ equals the sum of the main effect for S_{mh} , the main effect for B_{mh} , and their interaction, together with the constant implied by the omitted reference category. Thus, the

¹⁰The data for the figure are trimmed at the outdoor $\text{PM}_{2.5}$ level of $34.46 \mu\text{g}/\text{m}^3$, equivalent to the 99th percentile of the distribution of outdoor $\text{PM}_{2.5}$ values in weeks 9-12.

plotted values represent the level of indoor $\text{PM}_{2.5}$ implied by the estimated parameters, rather than any single regression coefficient.¹¹

Unlike [Figure 8](#), these estimates exploit within-household variation over time, eliminating concerns about differences across households exposed to different outdoor conditions. We find statistically significant decreases in indoor pollution over time at all except the lowest outdoor levels, with larger effects for higher outdoor pollution. At outdoor levels between 5 and 10 $\mu\text{g}/\text{m}^3$, indoor levels decline from 5.3 to 4.6 $\mu\text{g}/\text{m}^3$ between week 1 and weeks 9-12. At 30-35 $\mu\text{g}/\text{m}^3$ outdoor $\text{PM}_{2.5}$, the indoor reduction is substantially larger: from 20.4 to 14.2 $\mu\text{g}/\text{m}^3$. While outdoor pollution levels larger than 30 $\mu\text{g}/\text{m}^3$ are quite rare (3.6% of observations), large effects occur at more common, moderate outdoor pollution levels: at 20-25 $\mu\text{g}/\text{m}^3$ outdoor $\text{PM}_{2.5}$ (9.6% of observations show outdoor $\text{PM}_{2.5} > 20 \mu\text{g}/\text{m}^3$), indoor $\text{PM}_{2.5}$ falls from 14.7 at week 1 to 11.1 $\mu\text{g}/\text{m}^3$ at weeks 9-12.

This declining sensitivity suggests that households increasingly protect their indoor environments from outdoor pollution, consistent with defensive investments. This concentration of the largest effects at high outdoor pollution, where health damages are most severe, suggests households respond rationally to the risk information when it provides the greatest health benefits. Further, since we show that households manage to reduce their indoor pollution, given specific outdoor pollution, these results counter worries about mechanical mean reversion due to outdoor pollution reductions.

5 Household Mitigation Behaviors

Having established that IAQ improves following monitor installation, we now examine the specific actions households take. Recall that we use machine learning algorithms to infer two key behaviors from the pollution data itself: (1) reducing indoor pollution generation, particularly from cooking; and (2) installing air purifiers to filter indoor air. We also explore ventilation during cooking episodes for households without air purifiers by assessing how the evolution of IAQ following a cooking episode varies over time.

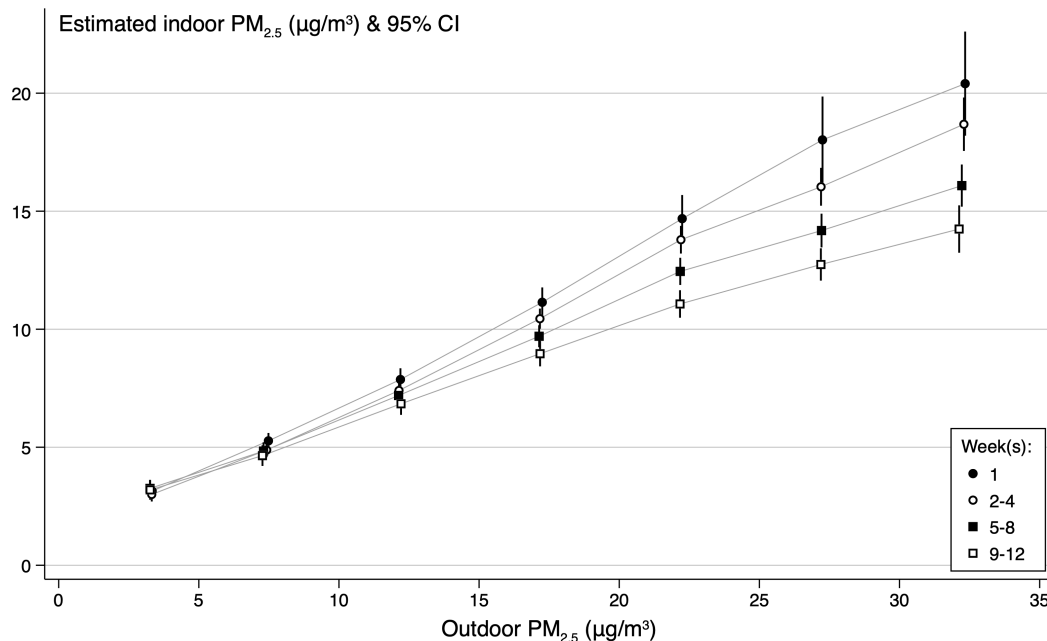
5.1 Changes in Cooking Behavior

[Figure 10](#) examines whether households alter their cooking behavior after learning about IAQ. Using our cooking detection algorithm to identify cooking episodes, we estimate [Equation 3](#): a Poisson pseudo-maximum likelihood regression where the dependent variable is a weekly count of cooking episodes. The figure presents incidence rate ratios for weeks 2-12 relative to week 1.

We find limited evidence that cooking frequency changes following monitor installation. The incidence rate ratios remain insignificantly different from 1 throughout the twelve weeks, except for the week-2 estimate, and point estimates range from 0.94 to 0.99. Cooking episodes show a statistically significant decrease in week 2, but revert to normal after. This stability suggests that

¹¹We follow the same approach in all other figures involving interaction terms, i.e., [Figure 11](#), [Figure 12](#), [Figure 13](#).

Figure 9: Estimated change in the indoor-outdoor pollution relationship by time since monitor installation



Estimated indoor PM_{2.5} (µg/m³) with 95% confidence intervals (based on joint estimates of the relevant parameters) from outdoor pollution bins 0-5, 5-10, 10-15, 15-20, 20-25, 25-30, and 30-35 µg/m³ for week 1 (filled circles), weeks 2-4 (hollow circles), weeks 5-8 (filled squares), and weeks 9-12 (hollow squares) after indoor monitor installation (α_s^2 , β_b^2 , and γ_{sb}^2 in Equation 2). The gray lines serve as a visual aid. **Notes:** The horizontal alignment of the estimates corresponds to the mean outdoor PM_{2.5} (x-axis) of the corresponding bin. To improve legibility, the estimates for the outdoor PM_{2.5} ≥ 35 µg/m³ bin are left out from the figure. The corresponding estimates and 95% confidence intervals are 33.85 [30.87, 36.84] for week 1, 34.29 [31.16, 37.41] for weeks 2-4, 29.80 [24.57, 35.04] for weeks 5-8, and 26.26 [20.80, 31.72] for weeks 9-12. The regression includes flexible controls for weather as well as indoor monitor, state-year-month, day-of-week, and hour-of-day fixed effects. Standard errors are clustered on Census tracts.

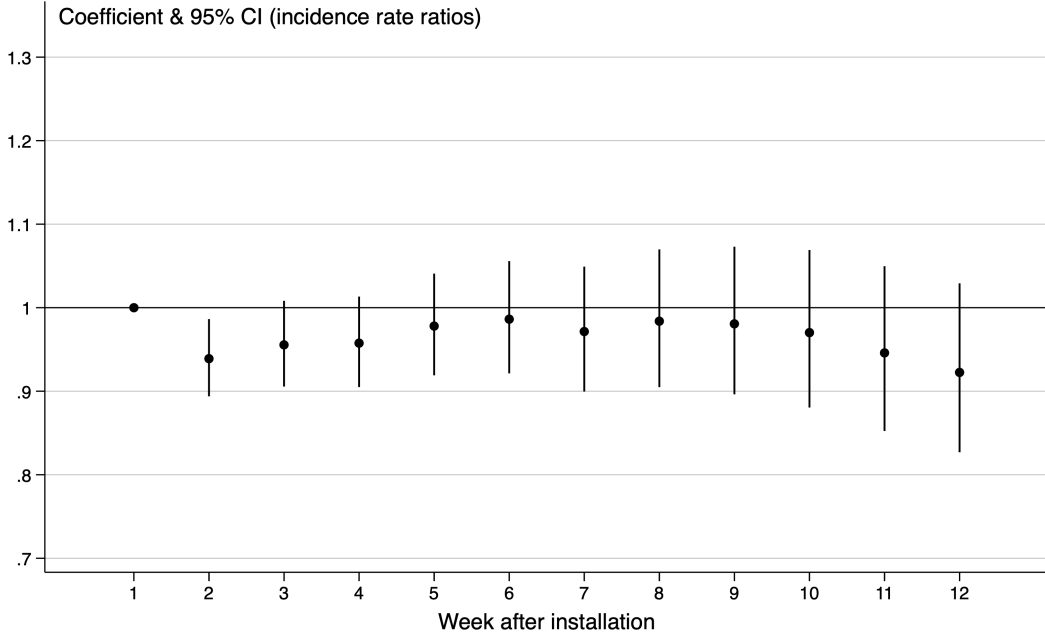
households do not reduce pollution generation as a mitigation strategy, perhaps because cooking is a necessary activity with costly substitutes.

5.2 Changes in Ventilation Behavior

While we cannot directly observe window opening or ventilation system use, we can infer ventilation changes from how cooking episodes affect IAQ (for households without air purifiers). Increased ventilation during or after cooking should reduce the magnitude and duration of cooking-related pollution spikes.

Figure 11 examines this by estimating Equation 4, which investigates how indoor PM_{2.5} evolves in the hours following a cooking episode onset, separately for bins of the n-th cooking episode: 1st (filled circles), 2nd or 3rd (hollow circles), 4th to 6th (filled squares), 7th to 12th (hollow squares) after monitor installation. Again, the plotted coefficients and confidence intervals are based on joint estimates of the relevant parameters. If households increase ventilation, we should see smaller or shorter-duration spikes in later cooking episodes.

Figure 10: Estimated change in cooking frequency



Coefficients with 95% confidence intervals, transformed to incidence rate ratios, from week after indoor monitor installation indicators (γ_w^3 in Equation 3). **Notes:** The regression sample only includes households with at least twelve recorded cooking episodes during the first twelve weeks after indoor monitor installation. The regression includes flexible controls for outdoor pollution and weather, as well as indoor monitor and state-year-month fixed effects. Standard errors are clustered on Census tracts.

The figure reveals stable IAQ patterns over time. Peak pollution levels during hours 1 and 2 after cooking onset remain constant at approximately $38\text{-}44 \mu\text{g}/\text{m}^3$ across all time periods. The decay rate back to baseline is also comparable over time. This stability suggests that households do not systematically alter ventilation practices during cooking, when ventilation would be most valuable for pollution reduction.

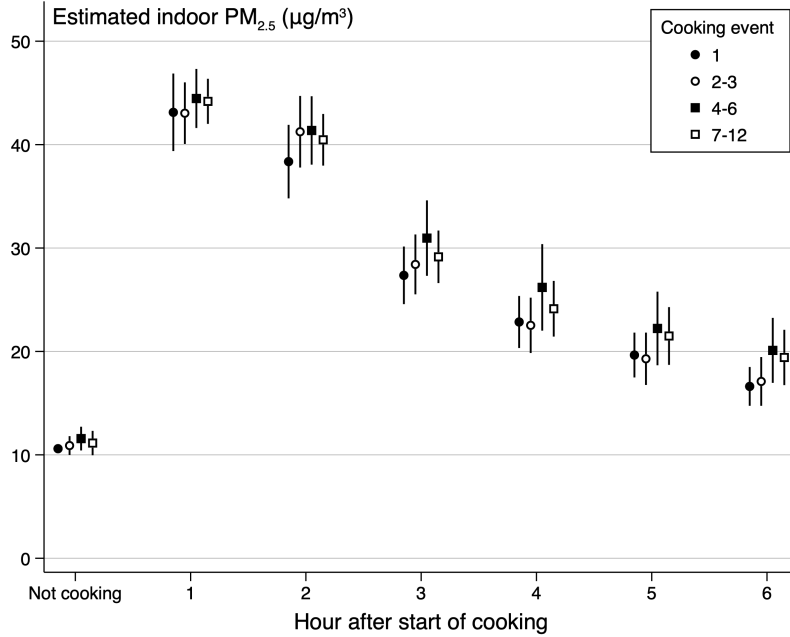
The absence of ventilation changes may reflect practical constraints. Opening windows during cooking can interfere with stovetop exhaust fans, affect indoor temperature control, and raise security concerns. These frictions may limit ventilation as a mitigation strategy despite its zero monetary cost. Another possibility is that households already had their windows open during cooking before installing the monitor due to preconceived notions that cooking generates indoor pollution.

5.3 Air Purifier Effectiveness

Our analysis suggests that air purifiers represent the primary household response to IAQ information. Figure 12 shows the results from estimating Equation 5, which examines the effectiveness of purifiers by comparing the indoor-outdoor gradient for households classified as non-adopters (filled circles) versus adopters (hollow circles).

Without purifiers, indoor pollution increases slightly less than one-for-one with outdoor levels.

Figure 11: Estimated change in ventilation during cooking




Estimated indoor PM_{2.5} (µg/m³) with 95% confidence intervals (based on joint estimates of the relevant parameters) for the 1st (filled circles), 2nd to 3rd (hollow circles), 4th to 6th (filled squares), and 7th to 12th (hollow squares) cooking episode after indoor monitor installation by hour after start of cooking with non-cooking times = 0 (α_n^4 , β_c^4 , and γ_{nc}^4 in Equation 4). **Notes:** The regression sample only includes households with at least twelve recorded cooking episodes during the first twelve weeks after indoor monitor installation, and before their first classification as running an air purifier. The regression includes flexible controls for outdoor pollution and weather as well as indoor monitor, state-year-month, day-of-week, and hour-of-day fixed effects. Standard errors are clustered on Census tracts.

With purifiers, indoor pollution remains substantially lower and increases more slowly with outdoor conditions. The gap between the two lines—the treatment effect of purifiers—grows from close to 0 µg/m³ at the lowest outdoor pollution bin to 11 µg/m³ at 30-35 µg/m³ and approximately 25 µg/m³ when outdoor pollution exceeds 35 µg/m³. These large effects during high pollution episodes, when health damages are most acute, suggest substantial health benefits from purifier adoption.

Figure 13 shows the results from estimating Equation 6, which assesses the impacts of air purifiers on cooking spikes (filled circles: non-adopters; hollow circles: adopters). We show the average IAQ outside of cooking episodes, which is largely a mechanical feature because we use this differential to classify households into purifier adoption. Since cooking episodes are excluded from the algorithm for detecting purifier adoption, the difference during cooking hours demonstrates the value of purifiers in improving pollution.

Outside of cooking episodes, PM_{2.5} is 2.4 µg/m³ lower once households use purifiers. Even during cooking, when indoor sources dominate, purifier adopters maintain lower pollution levels. While peak concentrations at the start of a cooking episode remain similar, the decay is much faster for adopters: the average difference between non-adopters and adopters is 7.2 µg/m³ during hours 2-6. Since our detection algorithm explicitly excludes cooking periods, these differences provide validation that we are capturing real purifier adoption rather than spurious patterns.

Figure 12: Estimated change in the indoor-outdoor pollution relationship by purifier status



Estimated indoor PM_{2.5} (µg/m³) with 95% confidence intervals (based on joint estimates of the relevant parameters) from outdoor pollution bins 0-5, 5-10, 10-15, 15-20, 20-25, 25-30, and 30-35 µg/m³ for households without air purifiers (filled circles) and households with air purifiers (hollow circles) (α^5 , β_b^5 , and γ_b^5 in Equation 5). Gray lines serve as a visual aid. **Notes:** The horizontal alignment of the estimates corresponds to the mean outdoor PM_{2.5} (x-axis) of the corresponding bin. To improve legibility, the estimates for the outdoor PM_{2.5} ≥ 35 µg/m³ bin are left out from the figure. The corresponding estimates and 95% confidence intervals are 38.56 [35.86, 41.25] without purifier, and 13.62 [12.71, 14.53] with purifier. The regression includes flexible controls for weather as well as indoor monitor, state-year-month, day-of-week, and hour-of-day fixed effects. Standard errors are clustered on Census tracts.


6 Conclusion

Our results demonstrate that private information about IAQ triggers substantial household mitigation efforts. Indoor PM_{2.5} concentrations fall by 2.5 µg/m³ on average within twelve weeks of monitor installation, with much larger reductions for high-pollution households. Upon installing monitors, households learn about their IAQ, which triggers defensive investments, primarily the adoption of air purifiers, rather than changes in pollution generation. The gradual nature of the response, with effects growing over twelve weeks, likely reflects the time households need to process monitor information and make purchasing decisions for durable goods like purifiers.

The concentration of effects among high-pollution households suggests rational responses to information. Households with good air quality show minimal response, avoiding unnecessary expenditures, while those with poor air quality invest in mitigation at rates highest where health benefits are greatest. This targeted response indicates that private information effectively channels defensive investments to where they generate the highest returns.

While several features support a causal interpretation of our results, we cannot fully rule out selection effects. We take several steps to limit concerns due to mean reversion and prior knowledge

Figure 13: Air purifier effectiveness during cooking



Estimated indoor PM_{2.5} (µg/m³) with 95% confidence intervals (based on joint estimates of the relevant parameters) for households without air purifiers (filled circles) and households with air purifiers (hollow circles) by hour after start of cooking with non-cooking times = 0 (α^6 , β_c^6 , and γ_c^6 in Equation 6). **Notes:** The regression sample only includes households with at least twelve recorded cooking episodes during the first twelve weeks after indoor monitor installation, and it only includes households with variation in air filtration. The regression includes flexible controls for outdoor pollution and weather as well as indoor monitor, state-year-month, day-of-week, and hour-of-day fixed effects. Standard errors are clustered on Census tracts.

of IAQ. However, the timing of monitor adoption may correlate with other unobserved factors that independently affect IAQ trajectories. If monitor purchases coincide with life changes that influence air quality, such as home renovations or changes in household composition, our estimates would conflate these effects with information responses. Despite these limitations, the patterns we document are consistent with households responding to new information about indoor air quality.

Since the households we study voluntarily purchase and install indoor air quality monitors, this select sample may lead to questions about external validity. These households may be more attentive to pollution risks or more responsive to information than the average household in the broader population. Our estimates should therefore be interpreted as characterizing the behavioral responses of individuals who actively seek personalized information, rather than as universal responses to information provision. Additional evidence from more representative samples would help clarify the extent to which these responses generalize beyond monitor adopters.

To gauge the welfare implications of these responses, we provide back-of-the-envelope calculations as a guide. Air purifiers cost \$100-300 initially with annual operating costs of \$50-100 for electricity and filters. Assuming reductions in PM_{2.5} persists for an entire year, an average of 2.5 members per household, and using estimates from Deryugina et al. (2019), the mean household in our sample experiences health benefits of \$725 annually from observed pollution reductions of 2.5

$\mu\text{g}/\text{m}^3$.¹² These benefits exceed the cost after only one year, with the gap likely to grow over time given the lower variable costs of air purifiers.

As sensor technologies continue to improve and costs decline, private risk information is likely to play an increasingly important role in environmental protection. Our results suggest this development could generate substantial welfare gains by enabling households to tailor defensive investments to their specific circumstances. At the same time, because publicly provided information is more equitable in its reach and lower-income households are less likely to adopt private monitors (Graff Zivin et al. 2025), the expansion of personalized monitoring introduces a potential efficiency–equity tradeoff. An important challenge for policymakers will be understanding how to ensure that accurate information reaches vulnerable populations who may benefit most from mitigation but face higher barriers to monitor adoption.

¹²This number might of course over- or understate benefits depending on our study population’s age and health profile.

References

- Aguilar-Gomez, Sandra (2025). Adaptation and mitigation of pollution: Evidence from air quality warnings. *Journal of Environmental Economics and Management* 134, 103215.
- Aguilar-Gomez, Sandra, Holt Dwyer, Joshua Graff Zivin, and Matthew Neidell (2022). This is air: The “nonhealth” effects of air pollution. *Annual Review of Resource Economics* 14(1), 403–425.
- Allcott, Hunt (2011). Rethinking real-time electricity pricing. *Resource and Energy Economics* 33(4), 820–842.
- Allen, Ryan W and Prabjit Barn (2020). Individual- and household-level interventions to reduce air pollution exposures and health risks: A review of the recent literature. *Current Environmental Health Reports* 7(4), 424–440.
- 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, Minwoo Hyun, and Jaecheol Lee (2022). Bounds, benefits, and bad air: Welfare impacts of pollution alerts. *NBER Working Paper 29637*.
- Barreca, Alan I, Matthew Neidell, and Nicholas J Sanders (2021). Long-run pollution exposure and mortality: Evidence from the Acid Rain Program. *Journal of Public Economics* 200, 104440.
- Barwick, Panle Jia, Shanjun Li, Liguang Lin, and Eric Yongchen Zou (2024). From fog to smog: The value of pollution information. *American Economic Review* 114(5), 1338–1381.
- Cheek, Emily, Valentina Guercio, Clive Shrubsole, and Sani Dimitroulopoulou (2021). Portable air purification: Review of impacts on indoor air quality and health. *Science of the Total Environment* 766, 142585.
- Chen, Chun and Bin Zhao (2011). Review of relationship between indoor and outdoor particles: I/O ratio, infiltration factor and penetration factor. *Atmospheric Environment* 45(2), 275–288.
- Chen, Yuyu, Avraham Ebenstein, Michael Greenstone, and Hongbin Li (2013). Evidence on the impact of sustained exposure to air pollution on life expectancy from China’s Huai River policy. *Proceedings of the National Academy of Sciences* 110(32), 12936–12941.
- Cheng, Kai-Chung, Hye-Kyung Park, Afua O Tetteh, Daisy Zheng, Nicholas T Ouellette, Kari C Nadeau, and Lynn M Hildemann (2016). Mixing and sink effects of air purifiers on indoor PM_{2.5} concentrations: A pilot study of eight residential homes in Fresno, California. *Aerosol Science and Technology* 50(8), 835–845.
- Chetty, Raj, Adam Looney, and Kory Kroft (2009). Salience and taxation: Theory and evidence. *American Economic Review* 99(4), 1145–1177.
- Coury, Michael, Liam Falconer, and Andrea La Nauze (2024). Wildfire smoke and private provision of public air-quality monitoring. *Journal of Environmental Economics and Management* 127, 103036.
- Deryugina, Tatyana, Garth Heutel, Nolan H Miller, David Molitor, and Julian Reif (2019). The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review* 109(12), 4178–4219.
- Deryugina, Tatyana and Julian Reif (2023). The long-run effect of air pollution on survival. *NBER Working Paper 31858*.
- Di, Qian, Yan Wang, Antonella Zanobetti, Yun Wang, Petros Koutrakis, Christine Choirat, Francesca Dominici, and Joel D Schwartz (2017). Air pollution and mortality in the Medicare population. *New England Journal of Medicine* 376(26), 2513–2522.
- Fan, Yichun (2024). Social cost of lifestyle adaptation: Air pollution and outdoor physical exercise. *Journal of Environmental Economics and Management* 127, 103042.
- Graff Zivin, Joshua, Benjamin Krebs, and Matthew Neidell (2025). Private adoption of public good technologies: The case of PurpleAir in California. *Environmental Research Communications* 7(6),

065025.

- Graff Zivin, Joshua and Matthew Neidell (2009). Days of haze: Environmental information disclosure and intertemporal avoidance behavior. *Journal of Environmental Economics and Management* 58(2), 119–128.
- Greenstone, Michael, Kenneth Lee, and Harshil Sahai (2021). Indoor air quality, information, and socioeconomic status: Evidence from Delhi. *AEA Papers and Proceedings* 111, 420–424.
- Handel, Benjamin R (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review* 103(7), 2643–2682.
- 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.
- Kelly, KE, J Whitaker, A Petty, C Widmer, A Dybwad, D Sleeth, R Martin, and A Butterfield (2017). Ambient and laboratory evaluation of a low-cost particulate matter sensor. *Environmental Pollution* 221, 491–500.
- Klepeis, Neil E, William C Nelson, Wayne R Ott, John P Robinson, Andy M Tsang, Paul Switzer, Joseph V Behar, Stephen C Hern, and William H Engelmann (2001). The National Human Activity Pattern Survey (NHAPS): A resource for assessing exposure to environmental pollutants. *Journal of Exposure Science & Environmental Epidemiology* 11(3), 231–252.
- 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.
- Lachowicz, Joanna Izabela, Simone Milia, Mariusz Jaremko, Enrico Oddone, Emanuele Cannizzaro, Luigi Cirrincione, Ginevra Malta, Marcello Campagna, and Luigi Isaia Lecca (2022). Cooking particulate matter: A systematic review on nanoparticle exposure in the indoor cooking environment. *Atmosphere* 14(1), 12.
- Liu, Fei-Ting, Kai-Ming Ting, and Zhi-Hua Zhou (2008). Isolation forest. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining (ICDM 2008)*, pp. 413–422. IEEE.
- Lloyd, Stuart (1982). Least squares quantization in PCM. *IEEE Transactions on Information Theory* 28(2), 129–137.
- MacQueen, J (1967). Multivariate observations. In *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability*, Volume 1, pp. 281–297.
- Metcalfe, Robert D and Sefi Roth (2025). Making the invisible visible: The impact of revealing indoor air pollution on behavior and welfare. *NBER Working Paper 33510*.
- Neidell, Matthew (2009). Information, avoidance behavior, and health: The effect of ozone on asthma hospitalizations. *Journal of Human Resources* 44(2), 450–478.
- Neidell, Matthew (2010). Air quality warnings and outdoor activities: Evidence from Southern California using a regression discontinuity design. *Journal of Epidemiology & Community Health* 64(10), 921–926.
- Noonan, Douglas S (2014). Smoggy with a chance of Altruism: The effects of ozone alerts on outdoor recreation and driving in Atlanta. *Policy Studies Journal* 42(1), 122–145.
- Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, 2825–2830.
- Pope III, C Arden, Majid Ezzati, and Douglas W Dockery (2009). Fine-particulate air pollution and life expectancy in the United States. *New England Journal of Medicine* 360(4), 376–386.
- Royer, Heather, Mark Stehr, and Justin Sydnor (2015). Incentives, commitments, and habit formation in exercise: Evidence from a field experiment with workers at a Fortune-500 company. *American Economic Journal: Applied Economics* 7(3), 51–84.
- Saberian, Soodeh, Anthony Heyes, and Nicholas Rivers (2017). Alerts work! Air quality warnings

- and cycling. *Resource and Energy Economics* 49, 165–185.
- Sater, Rita Abdel, Mathieu Perona, Elise Huillery, and Coralie Chevallier (2022). The effectiveness of personalised versus generic information in changing behaviour: Evidence from an indoor air quality experiment. *Unpublished Manuscript*.
- Shi, Xinjie, Yu Shen, and Ran Song (2023). Living with particles: Disclosure of pollution information, individual responses, and health consequences. *Journal of Health Economics* 92, 102824.
- Torkmahalleh, Mehdi Amouei, Soudabeh Gorjinezhad, Hediye Sumru Unluevcek, and Philip K Hopke (2017). Review of factors impacting emission/concentration of cooking generated particulate matter. *Science of the Total Environment* 586, 1046–1056.
- Ward, Alison L Sexton and Timothy KM Beatty (2016). Who responds to air quality alerts? *Environmental and Resource Economics* 65(2), 487–511.
- Xiang, Jianbang, Jiayuan Hao, Elena Austin, Jeff Shirai, and Edmund Seto (2021). Residential cooking-related PM_{2.5}: Spatial-temporal variations under various intervention scenarios. *Building and Environment* 201, 108002.
- Zhu, Yutong, Xiaoming Song, Rongshan Wu, Jiakun Fang, Lingyan Liu, Tong Wang, Shuo Liu, Hongbing Xu, and Wei Huang (2021). A review on reducing indoor particulate matter concentrations from personal-level air filtration intervention under real-world exposure situations. *Indoor Air* 31(6), 1707–1721.