Title
Essays in Environmental and Health Economics

Author
Wright, Catherine

Publication Date
2020

Peer reviewed|Thesis/dissertation

Essays in Environmental and Health Economics

by

Catherine Callahan Wright

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Meredith Fowlie, Chair
Professor Joseph Shapiro
Professor Jonathan Kolstad

Fall 2020

Essays in Environmental and Health Economics

Catherine Callahan Wright

Abstract

Essays in Environmental and Health Economics

by

Catherine Callahan Wright

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Meredith Fowlie, Chair

This dissertation contains three empirical studies that explore the health impacts of two dimensions of environmental quality, outdoor air pollution and extreme ambient temperature. All three studies rely on high-resolution, publicly available data on air pollution concentrations and weather, as well as restricted-access administrative data on individual-level mortality outcomes and hospital utilization. Chapters 1 and 3 explore the potential for health insurance to mitigate health damages from air pollution and temperature, respectively, while Chapter 2 incorporates recent advances in air particle dispersion modeling to assess whether health benefits projected under policies to reduce power plant emissions actually manifest in practice. The first two chapters generate causal estimates of the effects of poor air quality on health outcomes, while the third primarily provides suggestive evidence that public health insurance may play a role in driving heterogeneity in the relationship between extreme temperature and health.

In the first chapter, I investigate the potential for health insurance to mitigate adverse health impacts of air pollution in the United States – a country that lacks universal health care and continues to bear a large health burden from air pollution every year. To explore this issue, I focus on California as a case study, relying on exogenous variation in pollution exposure derived from wildfire smoke and a policy-induced increase in public health insurance coverage that varied in its regional intensity. Estimates for the state as a whole imply a 3% reduction in the health burden of air pollution for each one percentage point increase in the publicly insured rate. Back-of-the-envelope calculations based on these estimates imply a lower bound of $250 million per year in newly-measured benefits from California's ten percentage point Medicaid expansion under the Affordable Care Act. These results suggest that extending health insurance may be one way to address the increasing threat of wildfire smoke exposure in the western U.S., and potentially air pollution more broadly.

The second chapter, co-authored with Meredith Fowlie and Edward Rubin, focuses on the role of co-benefits in benefit-cost analyses of air pollution reduction policies. An environmental regulation generates pollution 'co-benefits' when it indirectly induces reductions in pollutants that are not the target of the authorizing legislation. In recent EPA rulemakings, quantified pollution co-benefits have exceeded directly targeted benefits. This increased prominence has invited increased scrutiny. We aim to advance the theory and inform the

practice of air pollution co-benefits accounting, focusing in particular on the health impacts of reduced exposure to small particulate matter ($PM_{2.5}$). We integrate a model of air particle dispersion into an empirical analysis of how reductions in power plant emissions impact local air quality downwind. We highlight the potential role of interactions with existing policies, documenting suggestive evidence that changes in power plant emissions have had differential impacts across areas that are more or less constrained by pre-existing air quality regulations. Future work will incorporate more precise modeling of pollution trajectories to explore these complex relationships, and associated health outcomes in greater detail.

In the third and final chapter, I return to the role of health insurance as a potential mediator of environmental health damages, in this case focusing on extreme temperature rather than air pollution. While the relationship between ambient temperature and human health has been studied at length by researchers across disciplines, this study is one of the first to consider the impacts of temperature on morbidity, rather than mortality, in a standard econometric framework. Relying on administrative hospital data from the state of Texas, I find that both extreme cold and extreme hot temperatures are associated with higher health care utilization, relative to a mild day. I also find that public health insurance may drive heterogeneity in the temperature-health relationship. Specifically, I estimate that health care utilization for temperature-related illness is relatively lower among adults insured by Medicaid, whereas the opposite holds for children. One potential explanation for this pattern is that the positive health benefits of public insurance strengthen over time, and that the lower relative benefit for children simply reflects their fewer number of potential years in the program. While these Medicaid-related estimates are suggestive rather than causal, this study provides novel evidence that income (which is correlated with Medicaid enrollment) may drive important heterogeneity in the relationship between climate and human health.

Acknowledgements

I would like to express my deepest gratitude to my advisor, Professor Meredith Fowlie, without whose expert guidance and constant support this dissertation would not have been possible. I would also like to thank Professor Ted Miguel whose mentorship since I was a Berkeley undergraduate was invaluable in setting me on (and keeping me on) this path. I am grateful to Professors Jon Kolstad and Ed Rubin, as well as my faculty and student colleagues at the Energy Institute, for their generosity in time and energy in helping me improve my research and hone my skills. I will always be grateful to my classmates, including Steph, Eric, Shelley, Stephen, Eva, Susanna, Karl, Steve, many others, and especially Danny Kannell, whose friendship and wisdom made the hardest parts of this experience more bearable. I also feel fortunate to have an uncle, Mark Perkins, whose efforts over the years to help strengthen my understanding of statistics have made me a better econometrician and researcher.

Dedication

I dedicate this dissertation to my loving and supportive parents, my stepmother, my sister, my grannie, and my entire extended family who have been a source of support and encouragement as I've ridden the highs and lows of this experience, and who have also suffered through many a dinner table conversation trying to help me understand my results. Finally, to Brittany, Shay, Janie, Karly, Eden, Jill, and the other Jeff: Thank you for always believing in me, celebrating the milestones, and cheering me on so loudly to the finish line.

Contents

Chapter 1

Does Health Insurance Reduce the Health Burden of Air Pollution? Evidence from California Wildfires

1.1 Introduction

Air pollution continues to pose a serious public health threat in many industrialized regions around the world, including in wealthy countries like the United States and Japan where current pollution levels are relatively low. The U.S. suffers nearly 100,000 premature deaths annually from exposure to particulate matter and ozone, and there is expected to be little improvement in this statistic over the next 30 years (OECD 2016). In the U.S., it is important that pollution-health impact estimates accurately reflect the true health burden of pollution exposure since these estimates are often the basis of national environmental policy, which in turn has important implications around the globe, including with respect to climate change.[1] More generally, efficient environmental policy requires accurate estimation of the magnitude of the externality, which may be mitigated (or exacerbated) by existing institutions and/or policies.

This paper assesses the interaction between air pollution externalities and U.S. health care policy. In places that lack a universal health care system, the propensity to invest in pharmaceuticals and other preventive or emergency care that could modify the costs of pollution exposure is governed in large part by access to health insurance. To study the influence of health insurance on the pollution-health relationship, this paper focuses on the case of California, a western state which has experienced an increasing wildfire (and associated smoke) burden in recent years and which simultaneously saw an over 40% expansion of its Medicaid insurance program beginning in 2014 under the Affordable Care Act. Leveraging spatial variation in exposure to the health insurance expansion and novel data on wildfire smoke plume trajectories, I find that enhanced access to public health insurance substantially mitigates the health burden of air pollution from wildfires. Relying on instrumented variation in local policy exposure, I estimate a local average treatment effect of approximately

[1] EPA frequently relies on research estimating the health impacts of pollution as the basis for benefit-cost analyses of new or modified environmental standards (e.g., U.S. EPA, 2012; U.S. EPA, 2011).

3% fewer emergency room visits and deaths related to smoky air for every one percentage point increase in Medicaid coverage. While the largest estimated impacts manifest in asthma-related ER visits and heart attack-related mortality, estimates of the relative effects of insurance and smoke are remarkably stable across the diagnoses and outcomes considered. I also find suggestive evidence that the Medicaid expansion has reduced the health burden from air pollution relatively more for Black and Hispanic residents than other demographic groups. Overall, assuming linearity, the estimates I obtain imply that for the approximately ten percentage point Medicaid expansion that California actually experienced as a result of the ACA, monetized health benefits related to wildfire smoke exposure are on the order hundreds of millions of dollars per year, at a minimum.[2] To my knowledge, this particular insurance benefit was not used to justify California's Medicaid expansion, and this is the first time that it has been measured empirically.

The findings of this paper contribute to a large empirical literature in environmental economics and related fields that measures the effects of air pollution exposure on mortality and other observable health outcomes, including emergency room (ER) utilization.[3] The earliest work in this area was primarily based on cross-sectional and longitudinal research designs, and damage estimates from this early work continue to inform some of EPA's long-standing environmental regulations. More recent work in this area has relied on sophisticated quasi-experimental methods to carefully identify the causal impacts of acute fluctuations in pollution exposure. For example, Schlenker and Walker 2016 exploit short-term variation in carbon monoxide (CO) exposure driven by congestion across U.S. airports to estimate that a one standard deviation increase in local CO levels is associated with an additional $540,000 in hospitalization costs for nearby residents. Of note, Schlenker and Walker measure adverse health effects at exposure levels below EPA's mandated maximum for the pollutant, suggesting possible room for re-calibration of current standards. In a related study of acute pollution fluctuations, Deryugina et al. 2019 estimate that moderate increases in local $PM_{2.5}$ concentrations driven by wind fluctuations are associated with non-trivial increases in elderly mortality. Miller, Molitor, and Zou 2017 estimate health impacts of a similar magnitude also among the U.S. elderly population, instead exploiting variation in $PM_{2.5}$ exposure driven by wildfire smoke, the same source of air quality variation that this paper leverages.

Beyond first-order health and health care impacts, pharmaceutical purchases by individuals are another way to value the (indirect) health costs of pollution, as demonstrated most clearly by Deschenes, Greenstone, and Shapiro 2017. Related work in the medical literature indicates that defensive actions – such as regular use of preventive pharmaceuticals – can yield significant health benefits, particularly for pollution-sensitive subpopulations. For example, several medical studies estimate an over 50% decrease in the relative risk of acute asthma exacerbations requiring emergency treatment among chronically asthmatic patients who are under long-term treatment with inhaled corticosteroids (see, for example, Donahue et al. 1997, Sin et al. 2004, and Merchant et al. 2018). In their seminal paper, economists

[2]The assumption that insurance benefits scale linearly is fairly strong. I discuss several caveats in Section 1.6.5.

[3]Researchers have studied the health effects of a variety of pollutants over the years, including ozone, carbon monoxide, and particulate matter of various sizes. Recent attention has often focused on fine particulate matter or "$PM_{2.5}$," generally considered to be the most harmful pollutant contained in wildfire smoke and thus the most relevant to this paper.

Deschenes et al. (2017) estimate that reductions in drug purchases during periods in which the *NO_x Budget Program* reduced ambient ozone levels by limiting industrial NO_x emissions were equivalent to over 60% of the monetized mortality benefits over the same period, or $800 billion dollars. Although Deschenes et al. acknowledge that "sticker" prices and co-pays for the pharmaceuticals they study often diverge significantly based upon on whether a patient is insured, the authors do not address at length the potential influence of health insurance coverage on the health outcomes they measure.

On the health insurance side, the conceptual starting point for this paper is that in the absence of a universal health care system, defensive (and emergent) health care utilization – both of which may influence the health impacts that researchers measure in response to pollution shocks – are price sensitive. Previous work on the economics of health insurance makes clear that prices strongly influence the type and amount of health care demanded by patients, in both randomized and observational evaluations. The RAND and Oregon Medicaid experiments found that lower out-of-pocket prices are associated with significant increases in the likelihood of taking prescription drugs, seeking outpatient care, and complying with recommended screenings (see, for example, Aron-Dine, Einav, and Finkelstein 2012 and Finkelstein et al. 2012). Quasi-experimental studies leveraging discontinuities in insurance eligibility for both public and private health plans have reached similar conclusions (e.g., Anderson, Dobkin, and Gross 2012, Card, Dobkin, and Maestas 2008, and Miller 2012). In each of these studies, the increased demand for health care documented by researchers primarily occurs in low-value care, a manifestation of the classic problem of insurance-related moral hazard and a phenomenon which policymakers have cited as a potential cost of expanding public health insurance.

The implications of insurance coverage for health care demand are slightly more nuanced in this paper's setting, in which air quality shocks derived from wildfire smoke exposure are both unexpected and temporal in nature. Conceptually, in the short-term, expanding health insurance involves countervailing effects on the demand for emergency care in response to poor air quality. A price effect predicts an increase in ER demand for a decrease in its price, while a "prevention effect" should act in the opposite direction and reduce ER visits, as newly insured individuals are more likely to consume preventive care, relative to previous levels and relative to the remaining uninsured population.[4,5] Over the longer run, health insurance may help enrollees adapt to the shock of smoke, further reducing ER visits without a corresponding increase in preventive care consumption, as insured individuals learn from previous smoke events and consume preventive care in earlier periods. Due to a narrow focus on the years surrounding California's 2014 Medicaid expansion as well as a lack of pharmaceutical data, the primary insights of this paper regard the short-run dynamics of insurance coverage for *net* ER visits and mortality. Of note, only the prevention effect is relevant for mortality outcomes, implying a weakly negative effect of health insurance on short-run mortality for pollution-related conditions when levels of pollution are high.

[4]Although preventive care (e.g., pharmaceuticals) in general represents low-value care, there are likely access/extensive margin constraints and/or liquidity constraints that make uninsured individuals less likely to consume such care to protect against poor air quality, even if it is relatively inexpensive.

[5]While health insurance almost certainly improves health outcomes for a wide variety of conditions, all else equal, the central idea of this paper is that insurance would disproportionately affect air quality-related health during periods when air quality is poor.

Building upon existing evidence that preventive care might play a substantal underlying role in the empirical estimation of pollution health impacts, as well as research showing that prices are strongly predictive of health care utilization, this paper is the first to consider health insurance as a mediator of short-run health damages from air pollution. The most closely related papers have studied potential health care-related mediators of another environmental shock, extreme temperature, but have generated mixed results (e.g., Barreca et al. 2016 estimate no impact of physician availability on heat-related health damages, while Mullins and White 2020 find that improvements in health care access following de-segregation of the American south mitigated some temperature-related illness, but only those related to extreme cold). Estimating the effect of health insurance on the pollution-health relationship will generate new knowledge about whether health insurance can be used to address the high health burden from air pollution that the U.S. continues to face. The results of this analysis may also inform the design of efficient environmental policy by improving the accuracy of impact estimates that underlie the calibration of regulatory thresholds.

The rest of the paper is organized as follows. Section 1.2 provides background information on the two key pieces of this analysis: wildfire trends in the western U.S. and the connection between smoke and $PM_{2.5}$, and the state of health insurance in California prior to and as a result of the ACA reform. Section 1.3 summarizes a simple model of demand for health care that elucidates why the expected short-run effect of health insurance on pollution-related ER visits is ambiguous. Sections 1.4 and 1.5 summarize the data and empirical strategy, while Section 1.6 reports results and provides discussion, including a counterfactual analysis based on empirical estimates. Section 1.7 concludes.

1.2 Background

1.2.1 Wildfire Smoke and Particulate Pollution

Wildfires are increasing in frequency and severity at a significant rate in the western United States. Since 1970, California has experienced an over 500% increase in annual area burned by wildfires, with particularly devastating wildfire seasons in 2017 and 2018 (Williams et al. 2019). In addition to the direct, immediate effects of wildfires in terms of loss of life and damage to wildlife and property, wildfire events can have severe indirect effects on human health through their creation of vast plumes of smoke. Further contributing to their perniciousness, wildfire smoke plumes can travel across large distances over relatively short time frames as dictated by unpredictable weather conditions, potentially resulting in exposure of wide swaths of the population living in areas far from the initial point of ignition.

While wildfire smoke contains many harmful chemicals, scientists generally agree that it is primarily the fine particulate matter or $PM_{2.5}$ which makes wildfire smoke so harmful to humans (Cascio 2019).[6] A recent study on the evolving burden of wildfire-derived air pollution in the U.S. confirms the strong empirical association between wildfire smoke and particulate pollution. Burke et al. 2020 link ground-level $PM_{2.5}$ concentration data to

[6]Human exposure to $PM_{2.5}$ can exacerbate underlying respiratory conditions such as asthma and chronic obstructive pulmonary disease (COPD), as well as trigger serious cardiovascular health events like stroke and heart attack.

satellite-based smoke plume locations and estimate that wildfires have accounted for between 25-50% of ambient $PM_{2.5}$ levels in recent years. Estimates of the statistical relationship between wildfire smoke exposure and $PM_{2.5}$ levels based on the California data used in this paper are reported in Table 1.1. In this table, *Smoke* is a binary variable that represents smoke coverage across ZIP-weeks or county-months of *at least* one standard deviation of the smoke coverage distributions for each of these aggregations. In each case, the one standard deviation threshold corresponds to full spatial coverage of the ZIP code or county by one or more smoke plumes on approximately one day during a week or four days during a month, respectively. The estimates in the table show a strongly positive correlation between smoke coverage and contemporaneous local $PM_{2.5}$ concentrations.

Table 1.1: Smoke Exposure and Ambient $PM_{2.5}$ in California (2012-2017)

	ZIP-Week		County-Month	
	Mean $PM_{2.5}$	Maximum $PM_{2.5}$	Mean $PM_{2.5}$	Maximum $PM_{2.5}$
	(1)	(2)	(3)	(4)
Smoke ($\geq$ 1 SD)	2.6***	5.2***	2.2***	13.9***
	(0.2)	(0.4)	(0.2)	(2.3)
Min. Temperature (°F)	−0.02	−0.1**	−0.01	−0.03
	(0.02)	(0.02)	(0.04)	(0.1)
Max. Temperature (°F)	0.1***	0.2***	0.04***	0.2***
	(0.01)	(0.02)	(0.01)	(0.1)
Precipitation (mm)	−0.03***	−0.03***	−0.4***	−0.6***
	(0.004)	(0.01)	(0.1)	(0.1)
Observations	44,926	44,926	3,710	3,710
Adjusted R^2	0.415	0.333	0.531	0.302

Notes: *p<0.1; **p<0.05; ***p<0.01. The dependent variables are the mean and maximum $PM_{2.5}$ concentrations in $\mu g/m^3$ based on official EPA monitors within a given ZIP-week or county-month. Not all ZIPs and counties in California have monitors. *Smoke* is a binary variable indicating smoke coverage of *at least* one standard deviation of the ZIP-week or county-month distributions, equal to 17.2% or 15.4% coverage, respectively. For conistency with the main specifications, all models include fixed effects for ZIP-week-of-year or county-month-of-year, and year. Standard errors are clustered by either ZIP or county.

The main source of air quality variation relied upon in this paper's empirical analysis is a measure of local smoke exposure. While a direct $PM_{2.5}$ elasticity would be easier to interpret and more externally valid than a reduced form smoke effect, estimates from the main empirical model of health outcomes are imprecise when instrumenting for $PM_{2.5}$ levels with smoke coverage. I therefore rely on smoke exposure itself as a proxy for local air pollution, noting that the estimates from Table 1.1 suggest that my preferred smoke threshold of one standard deviation or greater generally corresponds to an increase in local

$PM_{2.5}$ concentrations of approximately 2.5 $\mu g/m^3$.[7] Moreover, due to the unpredictable trajectories of plumes through space and time, exposure to smoke from wildfires appears to be as good as random, ensuring that estimates of the effect of smoke on health outcomes are unbiased. Further, wildfires and associated smoke exposure are increasingly policy-relevant and of growing concern to the public given wildfire trends in California and surrounding regions in recent years.

1.2.2 ACA in California

The Affordable Care Act was signed into federal law in 2010 before being fully enacted in 2014. The main components of the law were state-level Medicaid expansions to cover all adults with incomes below 138% of the federal poverty line, the creation of state-run private insurance marketplaces, expanded coverage for people under age 26 and those with pre-existing conditions, and the individual mandate. Taken together, these policy changes resulted in more than 20 million Americans gaining health insurance coverage over the first five years of the ACA, reducing the national uninsured rate from 15% to 10%. For data-related reasons, the focus of this paper is on California. Fortunately, the state is large, diverse, prone to smoke exposure from wildfires, and was significantly benefitted by the ACA. As a result of the health reform, California's overall uninsured rate fell by approximately 10 percentage points, from 17% in 2013 to 7% by 2016, the vast majority of which has been attributed to the increase in Medicaid enrollment of 6 million low-income Californians (KFF, 2020).[8]

Although the ACA substantially increased the insured rate among low income adults across California, the potential effect of the reform was largest in jurisdictions that were relatively *less* insured before the reform was enacted, an insight credited to Miller 2012 in her analysis of Massachusetts' 2006 health reform. This spatial variation in potential policy exposure is the main source of variation that I exploit in this paper to identify the causal effect of insurance coverage on health outcomes that respond to air pollution. Specifically, I augment a standard difference-in-differences (DID) research design by instrumenting for the continuous treatment variable of interest (the area-level Medicaid coverage rate) using plausibly exogenous pre-period variation in exposure to the policy. Figure 1.1 illustrates this variation in pre-reform insurance coverage at the ZIP code-level in California. For ZIPs with at least 10,000 residents, the 2013 adult uninsured rate ranged from 3% in 94507 (Alamo, Contra Costa County) to 64% in 90057 (MacArthur Park, Los Angeles County). At the county-level, this rate ranged from 13.2% in Marin to 37.7% in Alpine, a small county near Lake Tahoe.[9] Areas that were highly insured prior to the reform serve as natural controls and help to estimate the common time trend, while the remaining areas were treated by the policy to varying degrees based on the spatial differences depicted in Figure 1.1. Identification of the causal impacts of the reform could be threatened by differential trends in potential insurance coverage correlated with the timing of the Medicaid expansion, but I am not aware of any

[7] To put this number in context, the annual average and daily maximum thresholds that EPA enforces for $PM_{2.5}$ are set at 12.0 and 35.0 $\mu g/m^3$, respectively.

[8] For comparison, total enrollment in Covered California, the state's private insurance exchange created under the ACA, has held steady since 2014 at between 1.2-1.4 million individuals (PPIC, 2018).

[9] A map that corresponds to Figure 1.1 at the county-level is Appendix Figure A.3.

Figure 1.1: Adult Uninsured Rates, by ZIP Code (2013)

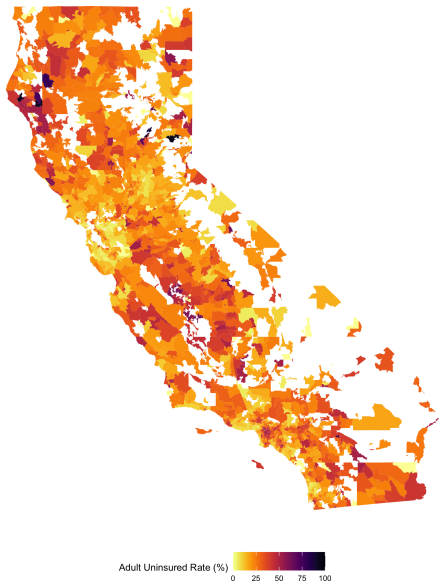

Notes: Figure shows the ZIP-level rate of uninsured adults in 2013. Data come from the American Community Survey 5-Year estimates. Blank areas of the map represents unincorporated areas that are not assigned a ZIP Code Tabulation Area (ZCTA), which is the spatial unit of analysis that corresponds to ZIP-based postal codes.

such trends. Still, to guard against the possibility of confounding, I control for observable socioeconomic indicators at the jurisdiction-level in each of the empirical specifications that follow.

As with any instrumental variables (IV) design, it is important to characterize the complier population, i.e. the individuals who are induced into treatment by the chosen instrument. In California, the great majority of new Medicaid enrollees post-ACA were low income adults between the ages of 18 and 65 and without dependent children.[10] One concern with a policy exposure instrument is that relevant demographics within this newly eligible group might vary systematically across high- versus low-exposure areas in ways that are correlated with outcomes, thus threatening the causal interpretation of empirical estimates.[11] More-

[10]This was the case for several reasons. Higher-income individuals and their dependents were generally covered by employer-based plans, which saw no change from the ACA, while low-income children and their guardians were often already eligible for existing public programs (e.g., Children's Health Insurance Program). Adults above the age of 65 are almost universally eligible for Medicare, the U.S.'s other major public insurance program.

[11]A related issue concerns the spatial overlap between areas highly exposed to the health care reform and the jurisdictions that are most affected by smoke from wildfires over the sample period. For example, causality (and external validity) could be threatened if the only areas that are exposed to smoke are those that are highly uninsured. To assess this, I compute the simple correlation between the fraction of smoke coverage and the pre-period adult uninsured rate in a given ZIP code or county. At the ZIP-level, the correlation ranges from 0.04% in the Inland Empire to 10% in the southern Central Valley. At the county-level, the range is from -6% in San Diego to 16% in the San Francisco Bay Area. These correlations appear low enough to rule out serious concerns about lack of overlap in the two primary treatment variables. In future work, however, it could be interesting to explore how pre-existing health care disparities might exacerbate

over, given that not all eligible individuals enroll in the programs for which they qualify, estimates will be further biased if the likelihood of program take up is correlated with treatment intensity. These concerns are especially important to address in this setting since I estimate potentially important race-based treatment effect heterogeneity for certain health outcomes.

I provide two pieces of evidence that speak to these concerns. First, Table 1.2 reports key ZIP-level demographics in the first year of the ACA across deciles of the pre-period adult uninsured rate (a proxy for treatment intensity). Unsurprisingly, as ZIPs move from less uninsured to more uninsured, per capita income falls. Simultaneously, while the fraction of Black residents remains relatively stable across uninsured deciles, the rate of Hispanic residents increases dramatically from about 10% to 50% when moving from the most insured to least insured decile. However, in the least insured areas, the proportion of white residents is still relatively high at about one third of total population. This breakdown provides some reassurance that, although Californias's Medicaid expansion disproportionately affected lower-income areas which tend to be less white and more Hispanic, a substantial number of the former do comprise a sizable fraction of total residents in the highest-exposure jurisdictions.

Table 1.2: ZIP Code Demographics, by Decile of Adult Uninsured Rate (2014)

Uninsured Decile	Uninsured Range (%)	Mean Income Per Capita ($)	Mean Population	White (%)	Black (%)	Hispanic (%)	Other Race (%)
1	0-7	49,093	6,714	70.0	3.4	11.6	15.0
2	7-11	50,293	20,005	66.3	2.7	12.5	18.5
3	12-15	43,025	23,435	61.8	3.6	17.2	17.4
4	15-18	37,633	24,168	61.5	3.6	18.8	16.1
5	18-22	31,904	22,595	58.9	4.2	23.2	13.6
6	22-25	28,434	23,465	57.0	4.0	27.5	11.6
7	25-29	25,404	23,153	49.2	4.4	34.4	12.0
8	29-34	22,741	24,948	42.3	6.1	39.3	12.3
9	34-41	19,809	23,405	35.0	4.3	50.8	9.9
10	41-100	17,137	21,429	32.4	4.4	56.1	7.1

Notes: Table reports unweighted statistics as of 2014, the first year of the ACA implementation, by decile of the adult uninsured rate in 2013. Unit of observation is the ZIP code. Other Race includes non-Hispanic individuals of any race not already represented in the table, or a combination of two or more races.

The second piece of circumstantial evidence on the characteristics of the complier group comes from the *California Health Interview Survey* in 2015, which describes differences between new enrollees and non-enrollees among the group of Californians that became newly eligible for Medicaid in 2014. The survey found that over 85% of newly Medicaid-eligible residents enrolled in the program in the first year, which is significantly higher than estimates of pre-ACA take up of between 50-60% (Wang and Trivedi 2017). Compared to newly eligible individuals that did not enroll, new enrollees were more likely to be female and over 35 years or older, and less likely to be employed. Although no significant differences were reported for rates of chronic asthma, new enrollees tended to have slightly poorer health,

environmental justice issues more broadly.

with elevated rates of obesity and heart disease (both of which could affect propensity to be harmed by exposure to air pollution). New enrollees also tended to have higher incomes (although all incomes in the sample fell below the new eligibility threshold of 138% of the federal poverty line). Most importantly for the heterogeneity analysis that follows, racial composition was not statistically different between enrollees and non-enrollees. This suggests that within the newly eligible population, racial differences did not drive differences in enrollment. Taken together, these statistics provide further reassurance that the differences I estimate across race stem from true underlying differences rather than statisical properties of my instrumented model (e.g., poor representation of certain demographic groups across all exposure levels, etc.).

1.3 Conceptual Model of Health Care Demand

This paper assesses how the effect of air pollution on health outcomes (ER visits and mortality) depends on access to health insurance. Intuitively, the expected direction of the effect of interest depends on the outcome considered. The expected short-run mortality impact from simultaneously improving health insurance access while reducing air quality is straightforward. An individual – whose likelihood of being insured is exogenously increased due to a change in policy – first chooses whether to seek preventive care (or take other related actions to avert or defend against pollution exposure). She subsequently experiences an unexpected, negative shock to local air quality, which may or may not result in her death. The only active choice in this scenario is whether or not to seek preventive care prior to the air quality shock. The likelihood of seeking preventive care will almost certainly increase after gaining insurance coverage, given that there exist both access constraints and liquidity constraints that make it relatively more cumbersome for the uninsured to obtain preventive care even if it is fairly inexpensive.[12] If the individual does seek preventive care, available medical evidence on the effectiveness of relevant pharmaceuticals suggests that a reduction in their price should have at least a *weakly negative* influence on the likelihood of death.[13] As noted previously, this protective effect should strengthen over time from the date of insurance enrollment, as new enrollees learn from temporal but seasonal smoke-related air quality shocks and become more likely to preemptively seek preventive care in anticipation of future shocks.

By contrast, the analysis of short-run morbidity outcomes (i.e., ER visits) is more nuanced. At a high level, insurance reduces the price of visiting the ER (which should increase ER visits, if they are a normal good), and also reduces the price of preventive and primary (i.e., doctor's office) care, which should decrease ER visits, if those types of care provide some sort of protection against acute illness episodes in response to environmental shocks. This protection effect may come through increased utilization of preventive care specifically to address the types of health conditions exacerbated by poor air quality, or it may come

[12]For context, the current price of budesonide, a common inhaled corticosteroid used for long-term asthma treatment, is currently about $300 without insurance (GoodRX.com) while the uninsured cost of an asthma-related ER visit is approximately $1,500 (Wang et al. 2014).

[13]There is no reasonable scenario under which increasing the probability of being insured would increase individual propensity to be fatally harmed by poor air quality, so the expected net effect of health insurance on pollution-related mortality is either null or negative.

through better management of chronic health conditions, thereby decreasing the likelihood that pollution exposure triggers acute illness. It is not clear *a priori* which of these countervailing effects – the price effect or the prevention effect – will dominate in the short run, which is the relevant time horizon in this paper.[14]

A simple conceptual model demonstrates mathematically why the expected direction of the effect of health insurance on air pollution-related morbidity is ambiguous in the short run.[15] This model is based on a highly simplified version of the canonical Becker-Grossman health production function model (Becker 1965; Grossman 1972), in which the production of health contributes directly and positively to individual utility. Specifically, consider a standard two-period utility maximization problem with no time discounting or uncertainty. The agent derives utility from health $H(\cdot)$, and from consumption of the numeraire good x, which has price normalized to unity. Health is a function of preventive care q, emergency (ER) care e, and pollution c. Investments in preventive care are made in $t = 1$, while the pollution shock and the choice of whether to seek emergency care occur in $t = 2$. Naturally, pollution is bad for health so $\partial H/\partial c < 0$, while preventive and emergency care contribute positively to health: $\partial H/\partial q > 0$ and $\partial H/\partial e > 0$. Insurance status is determined outside the model and sets the prices paid for both emergency and preventive care, p_e and p_q. This setup leads to a constrained utility maximization problem, where non-labor income Y is determined outside the system. The goal is to determine the sign of de/dc, which is how ER visits (e) are expected to change for a given increase in pollution levels (c), all else equal.

To determine the sign of de/dc, I proceed with a standard comparative statics analysis, the calculations for which are reported in Appendix Section A.2. These calculations show that the direction of this effect is indeterminate for at least three reasons. First, changes in relative health care prices (p_e and p_q) are unknown, and in particular, depend on the individual's baseline insurance status. It is likely that most uninsured individuals that became newly eligible for Medicaid through the ACA faced lower prices for all forms of health care post-reform, but that might not hold universally. Further, p_e and p_q are potentially correlated with e through the insurance choice, although this concern is mitigated if relying on an exogenous source of variation in coverage, as I do in the empirical analysis of this paper. Finally, the shape of the health production function $H(\cdot)$ is not known. In particular, it is not clear whether investments in q and e have decreasing, increasing, or constant marginal returns to health, or whether preventive and emergency care are compliments or substitutes. For these reasons, the theoretical direction of the effect of health insurance coverage on pollution-related ER visits is ambiguous, and thus should be assessed empirically.

[14] In the longer run, it is likely that the preventive effect will dominate, as health insurance coverage allows enrollees to learn from – and adapt to – repeated air quality shocks by seeking preventive care prior to the negative shock. In effect, insurance coverage will help individuals insure against poor air quality by making it easier and cheaper to consume low-value but highly protective care.

[15] Mullins and White 2020 present a similar conceptual framework in their paper on health care access and ambient temperatures, also taking care to distinguish between acute (emergency) and preventive care. The authors appear to make this distinction primarily in order to understand the possible mechanisms for their empirical results, since the health care expansion they study only affected access to preventive care. In contrast, the health care access shock that I study, the expansion of California's Medicaid program under the ACA, improved access to both preventive and acute health care simultaneously.

1.4 Data

1.4.1 Hospital Visits

Data on ER utilization comes from California's Office of Statewide Health Planning and Development (OSHPD). This database is a complete census of all patient visits that occurred at any hospital in the state over the period 2012-2017. Each OSHPD record represents a medical claim for a single visit, and contains demographic information about the patient, as well as details regarding their hospital visit, including diagnosis codes in order of relevance. To ensure a focus on visits that were likely induced by *current* air quality conditions, I filter the hospital sample to include only visits that originated in the Emergency Department (effectively excluding any pre-scheduled visits like planned surgeries, etc.). I link patients to contemporaneous, ambient air quality based on their home ZIP code and the week in which their hospital visit occurred.

I categorize ER patients into diagnosis groups based only on the first (primary) diagnosis code on the claim, in order to maximize statistical precision.[16] Diagnosis codes are based on the International Classification of Diseases (ICD), which is published by the World Health Organization (WHO) and used worldwide for morbidity and mortality statistics, insurance reimbursement systems, and other health care purposes. To identify diagnoses that are most likely to be affected by air pollution, I follow the convention in the pollution-health literature and first focus narrowly on asthma and heart attacks before expanding to include hypertension and visits for any respiratory-related condition. I also consider several "placebo" conditions – which should not be affected by air quality – in order to test the validity of my empirical model and specifications. For this, I follow the lead of previous researchers and examine effects of smoke and health insurance on ER visits for diabetes and appendicitis. The list of ICD codes used to identify each health condition are listed in Appendix Table A.1.

1.4.2 Mortality

To assess mortality impacts, I rely on restricted-access data on the universe of U.S.-based deaths contained in the Multiple Cause of Death (MCOD) files provided by the U.S. Centers for Disease Control and Prevention (CDC). The MCOD files are person-level records based upon individual death certificates for U.S. residents. To preserve individual privacy, each mortality event is identified only by the county and month in which it occured. Consequently, the mortality analyses that follow are conducted at the county-by-month level, whereas morbidity (ER visits) analyses are conducted at the ZIP-by-week level.[17] Each MCOD record contains a single underlying cause of death, up to twenty additional multiple causes, as well as demographic data about the decedent. The WHO defines the underying cause-of-death as "the disease or injury which initiated the train of events leading directly to death, or the circumstances of the accident or violence which produced the fatal injury." As with the hospital data, mortality data are linked to diagnoses using the ICD coding system,

[16] Analyses that consider secondary diagnoses generate similar, although weaker, regression estimates.

[17] To ensure comparability of estimates across these levels of aggregation, I also conduct morbidity analyses at the county-by-month level, which are reported in Appendix Table A.11.

and I assess mortality impacts on the same set of pollution-related and placebo conditions described previously.

1.4.3 Air Quality

The main source of air pollution variation in this paper is derived from the plausibly random trajectories of wildfire smoke plumes through space and time. Information on the location of smoke plumes in the atmosphere above the ground at a given point in time is a product of the National Oceanic and Atmospheric AdministrationâĂŹs Hazard Mapping System. This database contains daily geocoded outlines of smoke plumes that are generated through a combination of satellite imagery and visual screening by human analysts. The HMS smoke plume data cover the entirety of the contiguous U.S. and can be overlaid onto economically-meaningful spatial units, such as ZIP Code Tabulation Areas (ZCTAs)[18] or counties, using standard mapping tools.

Figure 1.2: Distribution of Wildfire Smoke Coverage (ZIP-Week)

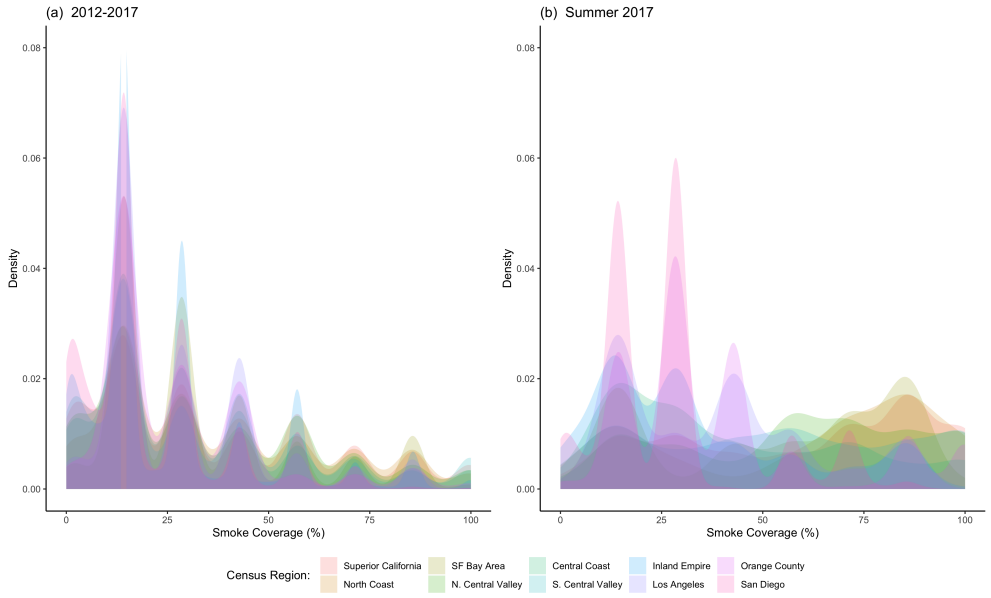

Notes: Figures plot the distribution of smoke coverage at the ZIP-week level in California, conditional on smoke coverage being greater than 0%. Smoke plume location data come from NOAA's Hazard Mapping System (HMS) and plumes have been overlaid onto ZIP Code Tabulation Areas (ZCTAs) from the U.S. Census Bureau. Panel (b) reports the distribution for Summer 2017 specifically because that was a severe wildfire year. Summer corresponds roughly to weeks 27-39 of the calendar year (June, July, and August).

I overlay smoke plumes on ZIP codes and counties, calculating the fraction of the spatial unit's area that was covered by smoke on a given day. I then aggregate the daily smoke coverage data to the ZIP-week (for the ER visits analysis) or the county-month (for the

[18]The connection between ZIP code and ZCTA is described in Appendix Section A.1.2. Throughout the paper, ZIP code refers interchangeably to both ZIP and ZCTA, since the differences between the two appear to be minimal for California.

mortality analysis) by taking the average coverage across the number of days in the time period. As an example, a smoke coverage measure of 50% could mean either that the geographic unit was fully covered by smoke plumes during half of the days in the period, or 50% covered by smoke plumes on every day of the period. For ease of interpretation, the main results presented in this paper all rely on a binary measure of smoke coverage, equal to a one standard deviation or greater level of smoke coverage across the full distribution of smoke. For the ZIP-week analysis, one standard deviation is 17.1% coverage (equal to 1.2 days of coverage out of seven days in a week) and for the county-month it is 15.4% (equal to 4.62 days of coverage out of 30 days in a month).

Figure 1.3: Distribution of Wildfire Smoke Coverage (County-Month)

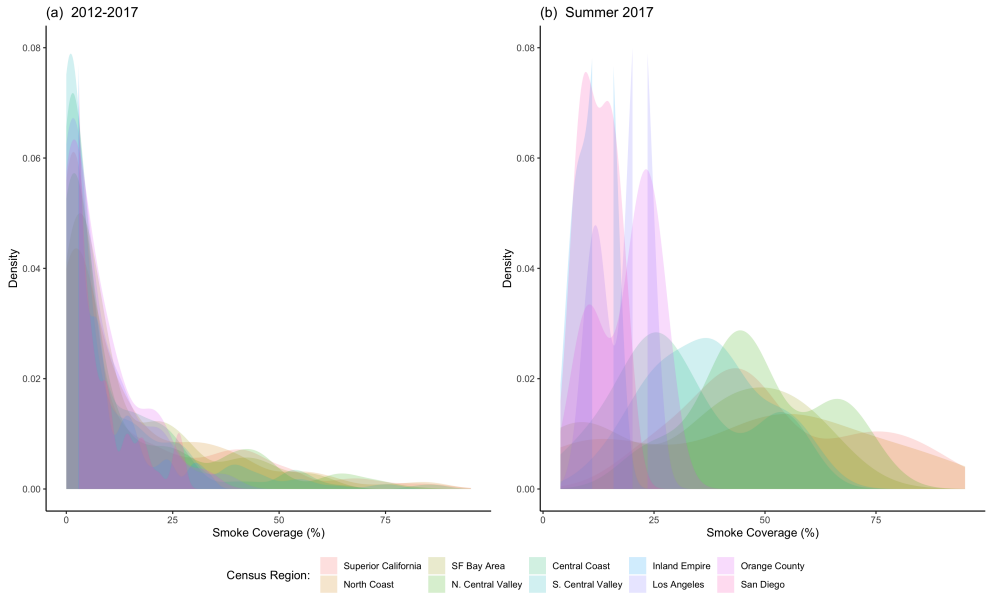

Notes: Figures plot the distribution of smoke coverage at the county-month level in California, conditional on smoke coverage being greater than 0%. Smoke plume location data come from NOAA's Hazard Mapping System (HMS) and plumes have been overlaid onto California counties. Panel (b) reports the distribution for Summer 2017 specifically because that was a severe wildfire year. Summer months are June, July, and August.

Figure 1.2 and Figure 1.3 show the distribution of smoke coverage for all years (Panel A) and limited to the summer of 2017 (Panel B) at the ZIP-week and county-month levels, respectively.[19] Comparing the two figures, counties are much less likely to be fully covered by smoke than are ZIP codes, which makes sense based on their relative sizes. The periodicity that is apparent in the ZIP-week figures corresponds to the total number of full-coverage days in a week. This is because the small geographic size of ZCTAs means that they are often fully covered by a smoke plume when they are covered by smoke at all. Additionally, the figures indicate that there exists significant variation in smoke coverage across regions above the one standard deviation threshold specified for the binary smoke indicator in the empirical specifications.

[19]The summer of 2017 was an especially smoky period, particularly in the San Francisco Bay Area and other parts of Northern California.

I supplement the smoke plume data with estimates of $PM_{2.5}$ ground-level concentrations, which are reported daily at the individual monitor-level in EPA's Air Quality System (AQS) database (along with monitor readings for all other criteria air pollutants that are regulated under the Clean Air Act). The network of official EPA monitors is relatively sparse, where some areas have just a single regulatory-grade monitor (e.g., San Francisco County) and other, often historically dirtier areas, have monitors in multiple locations (e.g., Los Angeles had 11 official EPA monitors for $PM_{2.5}$ as of 2018). I aggregate $PM_{2.5}$ monitor data to the ZIP-week and county-month by computing the mean and maximum readings across all available monitors in the area during a given week or month. Correlations between smoke coverage and local $PM_{2.5}$ concentrations are reported in Table 1.1.

1.4.4 Weather

I control for ambient maximum and minimum temperatures and total precipitation at the ZIP and county levels based on evidence of interactions between weather and air pollution. Data on daily temperatures and precipitation reported in a 4km-by-4km grid for the contiguous U.S. was obtained from PRISM Climate Group at Oregon State University. PRISM constructs measures for each grid cell by interpolating weather station data and accounting for relevant weather factors such as wind direction and topography.

1.4.5 Socioeconomic Variables

Socioeconomic data at the ZIP code and county levels comes from the American Community Survey (ACS) 5-year estimates (2008-2012 through 2013-2017). These variables are included as covariates in the main empirical specifications due to potential non-zero correlation of these measures with population health outcomes and/or insurance penetration. These data include area-level (ZIP code or county) measures such as the distribution of race, per capita income, unemployment, and health insurance status by age group (which is used to construct the instrument for the main IV-DID model). Monthly data on ZIP-level enrollment in California's Medicaid program (Medi-Cal) was obtained through a Freedom of Information Act (FOIA) request submitted to the California Department of Health Care Service (DHCS).

Table 1.3 reports summary statistics of all of the main variables at the county-month aggregation for the period 2012-2017, while the summary statistics for the same variables at the ZIP-week level are reported in Appendix Table A.2. As described, the measure of smoke coverage for the county-month aggregation represents the average spatial coverage during a single month in the given county. As is clear from the periodicity of the ZIP-week smoke distribution in Figure 1.2(a), ZIPs are often completely covered by smoke when they are covered at all. There is more variation in smoke coverage at the county-month level since counties in California are typically large enough to often be only partially covered by smoke plumes on a given day. Additionally, these summary statistics suggest that the majority of smoke exposure during 2012-2017 occured roughly in the summer months, or from June 1 through August 31.[20]

[20]In more recent years, wildfire events and smoke have continued into the late Fall months.

Table 1.3: Summary Statistics, 2012–2017 (County-Month)

Variable	N	Mean	p25	p75	Max
Morbidity outcomes:					
Total ER visits	4,176	19,171	1,766	19,721	325,841
ER visits: Asthma	4,176	315	24	292	8,718
ER visits: All Respiratory	4,176	1,992	177	1,989	62,137
ER visits: Hypertension	4,176	122	9	107	2,738
ER visits: Heart Attack	4,176	75	8	71	1,457
ER visits: Diabetes	4,176	175	14	160	3,178
ER visits: Appendicitis	4,176	53	3	51	1,015
Mortality outcomes:					
Total Mortality	4,176	362	32	384	6,229
Mortality: Asthma	4,176	18	2	18	347
Mortality: All Resp.	4,176	34	3	36	702
Mortality: Hypertension	4,176	11	0	11	220
Mortality: Heart Attack	4,176	7	0	8	215
Mortality: Diabetes	4,176	0	0	0	3
Air quality (Jun–Aug):					
% smoke coverage	1,044	27.0	9.4	40.9	95.1
Mean $PM_{2.5}$	929	9.4	6.5	11.2	106.8
Maximum $PM_{2.5}$	929	25.0	13.7	28.0	498.0
Air quality (Sep–Nov):					
% smoke coverage	1,044	2.3	0	1.7	37
Mean $PM_{2.5}$	928	10.0	5.6	12.5	55.0
Maximum $PM_{2.5}$	928	26.1	13.2	33.0	557.0
Air quality (Dec–Feb):					
% smoke coverage	1,044	0.3	0	0.2	10
Mean $PM_{2.5}$	928	9.1	5.0	11.0	50.0
Maximum $PM_{2.5}$	928	22.7	11.2	30.2	355.4
Air quality (Mar–May):					
% smoke coverage	1,044	3.3	0	4.8	26
Mean $PM_{2.5}$	925	7.1	5.1	8.7	42.7
Maximum $PM_{2.5}$	925	15.7	10.5	18.6	76.5
Population and health insurance:					
Total population	4,176	649,465	45,173	692,696	10,032,014
Adult population	4,176	410,040	26,564	417,688	6,493,977
% Uninsured adults, 2013 (Z)	4,176	23.2	18.5	27.6	37.7
% Medicaid coverage	4,176	29.4	21.6	36.8	56.6

Notes: Table reports unweighted statistics for the main mortality estimation sample. Unit of observation is the county-month. $PM_{2.5}$ data is only available for the counties that have official EPA monitors.

1.5 Empirical Strategy

An empirical test for the primary research question of this paper is to estimate the joint effect, on ER utilization and mortality separately, of short-run exposure to air pollution (derived from wildfire smoke) *and* improved access to public health insurance. I model these two distinct relationships using the same general regression framework:

$$\begin{aligned} log(Count_{jty}) = \alpha\ Smoke_{jty} + \beta\ MedicaidRate_{jty} + \\ \tau\ (Smoke_{jty} \times MedicaidRate_{jty}) + \\ X'_{jty}\ \gamma + \theta_{jt} + \delta_y + \varepsilon_{jty}, \end{aligned} \tag{1.1}$$

where Y_{jty} is the log of the count of either ER visits or deaths, for a given health condition, in jurisdiction j over time period t in year y. $Smoke_{jty}$ is a measure of smoke plume coverage over jurisdiction j during time period t,[21] and $MedicaidRate_{jty}$ measures the rate of Medicaid coverage in jty (in percentage points of the local population). The parameter of primary interest is τ, the coefficient on the interaction between smoke and Medicaid coverage, which measures the percent change in smoke-related ER visits or deaths associated with a one percentage point increase in the rate of public health insurance when the air is smoky.[22]

The coefficient on the Medicaid rate, β, measures the effect of public health insurance on morbidity and mortality rates, and it is important to note that its interpretation is different depending on the outcome considered. For ER visits, this effect represents a general "price effect" of insurance when the level of smoke held constant. If ER utilization is a normal good, basic economic theory suggests that a Medicaid program expansion (which increases the disposable income of eligible households that enroll and no longer have to pay for private health insurance and/or costly medical services if uninsured) should be associated with an increase in quantity consumed, in the absence of an adequate subsitute. However, enrolling in Medicaid simultaneously involves price reductions of other forms of health care, including primary care (e.g., office visits) and prescription medications, both of which new enrollees may substitute toward when becoming insured. As discussed in Section 1.3, shifting toward these health care substitutes could in theory reduce the quanty of acute ER care demanded. Therefore, the main effect of health insurance on ER demand in response to smoke is ambiguous. This result is shown mathematically in Appendix Section A.2. By contrast, the effect of insurance on mortality should be unambiguously (weakly) negative.

The dependent variable in Equation (1.1) represents a week-based total (or a month-based total, depending on the outcome measured). This temporal aggregation should somewhat alleviate concerns about "harvesting" (i.e., temporal displacement of health outcomes) which is a standard concern in studies that rely on daily or even shorter-term variation in environmental quality. The variables included in X_{jty} represent a selection of weather-related and socioeconomic controls that may independently affect the rates of ER utilization and mor-

[21] For the main specifications in the paper, this is constructed as a binary variable representing smoke coverage greater than or equal to standard deviation of the distribution of smoke coverage for ZIP-weeks or county-months.

[22] Under the assumption that the underlying population counts trend smoothly, the estimated discontinuities in log ER visit and death counts that I obtain can be interpreted as estimates of the percentage discontinuities in rates (Card, Dobkin, and Maestas 2004 provide a justification of this approach).

tality. These include temperature distributions, total precipitation, the fraction of j that is non-white, and median per capita income. Equation (1.1) also includes ZIP by week-of-year or county by month-of-year (θ_{jt}) and year (δ_y) fixed effects. The ZIP-by-week-of-year (or county-by-month-of-year) fixed effects represented in θ_{jt} control for any seasonal correlation between smoke and health outcomes that is common across years in given jurisdiction. δ_y is a time trend that controls flexibly for common time-varying shocks, such as those that induced by any changes to health care or environmental policies that apply universally to all jurisdictions in California. Identifying variation comes from differences across years (2012-2017) within ZIP-week-of-year or county-month-of-year. I estimate alternative specifications with varying controls and fixed effects to demonstrate the robustness of my results. For the morbidity and mortality analyses, standard errors are clustered at the level of the ZIP code and county levels, respectively. This clustering allows for arbitrary within-jurisdiction correlation.

After simplifying the air quality dimension of the model by defining $Smoke_{jty}$ to be a binary variable representing smoke coverage of at least one standard deviation, Equation (1.1) can be viewed as a standard difference-in-differences (DID) model with a continuous treatment measure, equal to the level of Medicaid coverage in jt. This treatment measure changed discontinuously in 2014 as a result of California's expansion of its Medicaid program under the ACA. However, if estimated by OLS, coefficients in Equation (1.1) are likely to be biased because of underlying (almost always unobservable) relationships between health insurance coverage, socioeconomic status, baseline health, risk preferences and other factors that influence the propensity to seek medical care, or to become seriously ill or die, when air quality worsens. For example, lower income individuals are more likely to be uninsured and in poorer health, which might both affect susceptibility to poor air quality. A credible strategy for causal inference requires isolating exogenous variation in Medicaid coverage that is uncorrelated with unobservable propensities for acute and/or fatal illness during periods of high air pollution. I therefore employ an IV strategy, using the pre-period (2013) adult uninsured rate in jurisdiction j as an instrument for the Medicaid rate in jt in the post-period. Specifically, I define $Z_{jy} = PrePeriodAdultUninsuredRate_{j,2013} \times Post_y$ where $Post_y$ is an indicator for years in the sample period after the policy was implemented (e.g., 2014-2017). The empirical specification for the first stage is then

$$MedicaidRate_{jty} = \lambda\ Smoke_{jty}\ +\ \rho\ Z_{jy}\ +\ X'_{jty}\ \gamma + \theta_{jt} + \delta_y\ + \xi_{jty}, \qquad (1.2)$$

where Equation (1.2) contains the same set of covariates and controls that are included in the original endogenous model. I then estimate the following equation using the instrumented treatment variable $\widehat{MedicaidRate}_{jty}$ that contains only the variation in jurisdiction-level Medicaid coverage that is explained by the pre-period uninsured rate, a proxy for exposure to the policy change.

The second stage equation, Equation 1.3, is an IV-DID intensity model, where τ_{IV} estimates a local average treatment effect (LATE) of a one percentage point increase Medicaid coverage, during periods of smoke, on health outcomes for new Medicaid enrollees living in areas that were highly uninsured prior to the implementation of the ACA in 2014. In this research design, areas that had complete (or nearly complete) health insurance coverage in 2013 serve as the pure control group which was unaffected by the policy change. If access

to health insurance offers some sort of protective value during periods of poor air quality – through enhanced access to preventive care, or something else – ER utilization (and mortality) should fall in areas that were more exposed to the policy as compared to areas which had nearly full insurance pre-reform. Therefore, we would expect $\widehat{\tau}_{IV}$ to be negative. The second stage IV-DID model is

$$\begin{aligned} log(Count_{jty}) = \alpha\ Smoke_{jty} + \beta\ \widehat{MedicaidRate}_{jty} + \\ \tau_{IV}\ (Smoke_{jty} \times \widehat{MedicaidRate}_{jty}) + \\ X'_{jty}\ \gamma + \theta_{jt} + \delta_y + \varepsilon_{jty} \end{aligned} \tag{1.3}$$

There are several conditions that must be met for this empirical model to identify a causal effect. First, identification relies on the assumption that if the ACA had not been enacted, deviations in ER utilization and deaths from their average levels in high- and low-uninsurance areas would have evolved similarly. This is the parallel trends assumption for difference-in-differences models that include place-by-temporal-unit fixed effects (i.e., θ_{jt}). I proceed by evaluating whether pre-reform trends in these outcomes were similar across areas with differing uninsured rates prior to the reform. The intuition behind this sort of placebo-based test is to check for statistically identifiable treatment effects on relevant outcomes in the years before the new policy was in place. Specifically, I estimate a model similar to Equation (1.3) in which instead of interacting the 2013 adult uninsured rate with an indicator for post-policy years, I interact it with individual indicators for each year in the 2012-2017 sample period. The excluded reference year is 2013, which leaves the 2012 interaction for assessment of pre-policy trends. Reassuringly, for ER visits and deaths associated with almost all of the outcomes considered, the placebo coefficients are indistinguishable from zero. The only detectable pre-trends are observed for all respiratory conditions (ER visits and mortality) and diabetes (mortality only), although these health conditions are not the main focus of this paper. I report pre-trend coefficients for all of the diagnoses and each health outcomes in Appendix Section A.5.1. These placebo tests provide suggestive evidence that the relevant outcomes were not trending at differential rates across areas with high versus low insurance in pre-reform years, and thus that any post-reform changes are attributable to the change in policy.

Another concern with identification is that difference-in-differences models will produce biased treatment effect estimates if there are time-varying, place-specific shocks that are correlated with outcomes of interest and occur at the same time as treatment. This type of confounding results in a violation of the parallel trends assumption, but it is unfortunately not testable with a similar placebo-based test due to non-zero correlation with the timing of treatment itself. In this paper's setting, a violation of this condition requires that something else changed in aggregate in 2014 that could influence health outcomes related to smoke exposure, but was experienced by different jurisdictions to varying degrees. For instance, if areas with relatively low pre-ACA Medicaid penetration rates experienced jurisdiction-level shocks that made it easier to access public health care starting in 2014 (for reasons unrelated to the Medicaid expansion), then this effect could confound the effect of the public insurance expansion that the focus of this paper. However, the possibility of confounding at this level seems unlikely in this setting. The most significant health care-related, aggregate

changes that occurred at the same time as California's Medicaid expansion were the other components of the Affordable Care Act, and each of those policy changes applied equally across all jurisdictions in the state (and to many extents, the U.S. as a whole). Additionally, differential socioeconomic trends across jurisdictions that could affect health outcomes – such as changes in employment or poverty rates – are observable and included in X_{jty} in Equation (1.3), providing further defense against potential confounding of the estimated treatment effect.

A third assumption that must be met for this model to deliver unbiased, causally-identified estimates is related to the instrumental variables component: relevance of the instrument, Z_{jy}. Relevance requires that there be sufficient correlation between the instrument and exogenous regressor, since "weak" instruments can produce IV coefficient estimates that perform poorly even relative to OLS (for reasons surveyed in Stock, Yogo, and Wright 2002). The standard test for weak instruments is an F-test on the first stage regression. The F-statistic for the model estimated in Equation 1.2 is 55.8, far exceeding the standard threshold of 10, and indicating that the instrument is highly relevant to the endogenous ZIP-level Medicaid rate. This positive correlation is depicted graphically in Figure 1.4, which plots the unconditional distribution of changes in ZIP-level Medicaid penetration from 2013 to 2017 across deciles of the adult uninsured rate in 2013. While there is variation in the endogenous variable within each decile, there is a clear monotonic relationship between uninsured rates and mean Medicaid growth. The county-month version of Figure 1.4 is Appendix Figure A.2, and first stage estimates at the county-level are reported in column (4) of Table A.5). The F-statistic for the county-level first stage equation is 265.9.

Figure 1.4: First Stage Correlation (ZIP-Level)

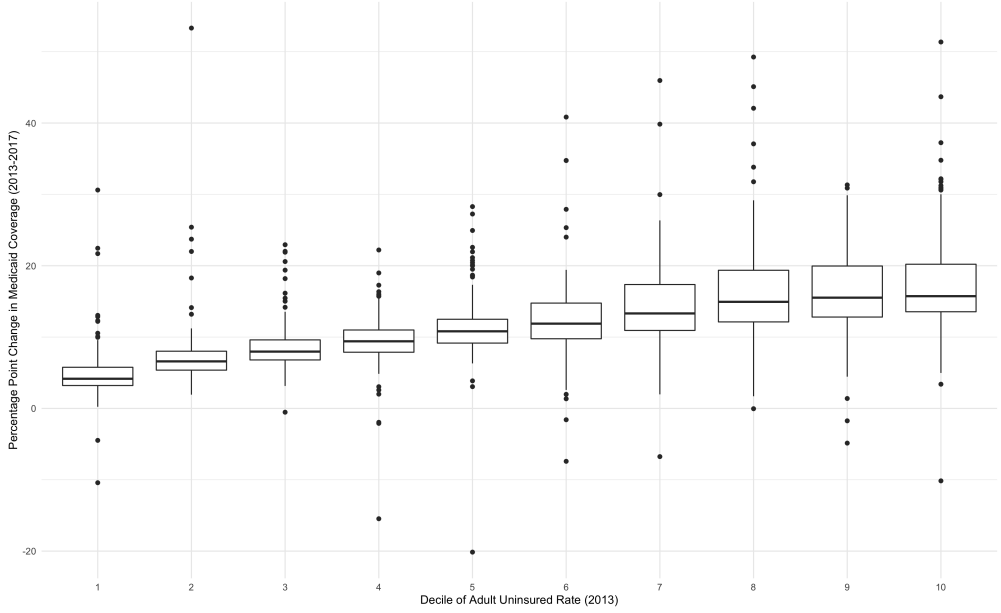

Notes: Figure shows the unconditional distributions of the percentage point change in ZIP-level Medicaid coverage from 2013 to 2017 for each decile of the exogenous instrument (the adult uninsured rate in 2013). ZIP codes with total population smaller than 500 (134 out of 1619 total with non-missing insurance data) have been dropped.

The final condition that must be satisfied for causal identification of τ_{IV} is the exclusion restriction. This requires that the instrument not have a *direct* effect on the outcome(s) of interest (ER visits and mortality), implying that the instrument is safely excludable from the original model in Equation (1.1). The exclusion restriction further requires that the instrument not have an *indirect* effect on outcomes through non-zero correlation with any relevant omitted variables. This is equivalent to requiring that $Cov[Z_{jy}, \varepsilon_{jty}] = 0$, or that the instrument is not correlated with the error term in the baseline model, since this would violate the standard exogeneity requirement. Excludability and exogeneity of the instrument in general are not conditions that can be tested empirically. However, after controlling for observable time-varying socioeconomic measures that in theory or expectation could influence health outcomes it seems unlikely that a single-year, jurisdiction-specific uninsured rate would be systematically related underlying propensity to use the ER or die when exposed to air pollution from wildfire smoke plumes.

1.6 Results

1.6.1 ER Visits

As discussed, in the results that follow, $Smoke_{jty}$ is constructed as a binary variable representing smoke coverage in ZIP j and week t of year y of *at least* one standard deviation of the distribution of smoke coverage across ZIP-weeks during 2012-2017 in California. For the analysis of ER visits, this is equal to 17.2% coverage or full coverage during approximately 1.2 days in a week. This modeling choice is intended to simplify the interpretation of magnitudes in the reported estimates by forcing the intensity effect of the main IV-DID intensity model to come solely through changes in the Medicaid rate (and not simultaneously through changes in the intensity of smoke exposure). Recognizing that this represents a non-trivial simplification, I also conduct a supplemental analysis in which I instrument for $PM_{2.5}$ using the fraction of smoke coverage in a jurisdiction (which I refer to as the "double-IV-DID" model). First-stage estimates are reported in Table 1.1, and indicate that the average increase in $PM_{2.5}$ in a given ZIP-week or county-month associated with this binary smoke variable is about 2.5 $\mu g/m^3$. Reassuringly, the second-stage estimates generated by this double-IV model in general are not meaningfully different in magnitude from the results contained in the main text. However, restricting the sample to only counties and ZIP codes that contain official EPA $PM_{2.5}$ monitors results in very few precisely measured estimates, and therefore these results are reported in Appendix Section A.6, rather than the main text.

Table 1.4 reports OLS, first stage, reduced form, and IV estimates of the relationship between smoke and health insurance on asthma-related ER visits. Medicaid coverage rates and the 2013 adult uninsured rate (Z) have been scaled to represent an increase in each of ten percentage points. This scaling was implemented since it represents approximately the reduction in California's uninsured rate as a result of the ACA (the vast majority of which was due to the Medicaid expansion), and it is also approximately the remaining uninsured rate as of 2020. Column (1) of Table 1.4 reports baseline OLS estimates that only include spatial and temporal fixed effects, while column (2) includes additional weather and socioeconomic controls to reduce bias and increase statistical precision. Columns (3)-

(6) report corresponding first stage, reduced form, and IV estimates of the causal effect of smoke exposure and health insurance when instrumenting for the jurisdictional Medicaid rate with Z (the adult uninsured rate in 2013). The dependent variable in the first stage equation, column (4), is the rate of Medicaid coverage in jty in percentage points of the local population. For all other columns, the dependent variable is the log of the count of asthma-related ER visits.

The estimates in Table 1.4 indicate that Medicaid insurance reduces the health damages of air pollution. Preferred IV estimates are in column (6), and are substantially larger than corresponding OLS estimates in column (2), suggesting that OLS suffers from attenuation bias in this setting. It is first worth noting that the OLS and IV estimates capture the behavior of different populations, which may have very different health care demand responses to both air quality and health insurance. While the OLS estimate is an average treatment effect for the population as a whole, the IV estimate is a local average treatment effect for the population induced into treatment by the instrument (e.g., new enrollees living in areas that were highly uninsured pre-ACA). The attenuated Medicaid OLS estimate suggests that compliers use the ER relatively more for asthma-related issues than does the overall population, when holding smoke constant. Conversely, compliers are less likely than the average person to visit the ER for asthma issues when smoke crosses the one standard deviation threshold (represented by the negative coefficient on the main interaction between smoke and Medicaid). These results together suggest there might be some important differences between new and existing Medicaid recipients that lead the former to seek more emergency health care in general (as evidenced by the positive Medicaid coefficient) but seek it less in response to poor air quality (with a negative interaction coefficient possibly indicative of a larger prevention effect among new enrollees). These differences between OLS and IV also indicate caution is warranted if intending to extrapolate these IV estimates to the overall population.

Table 1.4: OLS and IV Estimates — ER Visits (Asthma)

	OLS	OLS	First Stage	Reduced Form	IV	IV
	(1)	(2)	(3)	(4)	(5)	(6)
Smoke	0.012***	0.010**	−0.020***	0.020***	0.008***	0.046***
	(0.004)	(0.005)	(0.005)	(0.004)	(0.003)	(0.007)
Medicaid	0.004	0.005*			0.020**	0.023**
	(0.003)	(0.003)			(0.010)	(0.010)
Smoke × Medicaid	−0.001	−0.001				−0.013***
	(0.001)	(0.001)				(0.002)
Z			0.228***	0.006***		
			(0.026)	(0.002)		
Smoke × Z				−0.008***		
				(0.001)		
Included controls:						
Year FE	✓	✓	✓	✓	✓	✓
ZIP-week-of-year FE	✓	✓	✓	✓	✓	✓
Weather controls		✓	✓	✓	✓	✓
Socioeconomic controls		✓	✓	✓	✓	✓
Observations	551,019	543,052	543,052	545,635	543,052	543,052
Adjusted R^2	0.774	0.771	0.902	0.772	0.771	0.771

Note: *p<0.1; **p<0.05; ***p<0.01. The dependent variable in columns (1)-(2) and (4)-(6) is the log of the count of asthma ER admissions in ZIP j and week t of year y. The dependent variable in column (3) is the Medicaid coverage rate (the total number of Medicaid enrollees divided by total population) in ZIP j and week t of year y. Socioeconomic controls include fraction Black, fraction Hispanic, fraction of other race, overall population, and per capita income. Medicaid-related estimates have been scaled to reflect a ten percentage point increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across ZIP-weeks (equal to 17.2% coverage) or greater. Standard errors are clustered by ZIP.

The IV estimates in Table 1.4 imply that smoke coverage of 17.2% or more during a single week causes a 4.6% increase in asthma-related ER visits by local residents, an effect of which over 25% (i.e., 1.3%) is mitigated by raising the Medicaid coverage rate by ten percentage points. Further, the ER price effect (captured in the coefficient on the main Medicaid effect) of 2.3% represents the overall increase in ER visits for asthma holding smoke-related air pollution at less than 17.2% coverage per week. A negative coefficient on the interaction term provides suggestive evidence that the prevention effect dominates the price effect when ambient air is relatively smoky. The first stage estimates in column (3) imply that a one percentage point increase in the adult uninsured rate in 2013 (Z) is associated with a 0.23 percentage point increase in Medicaid coverage in the post-period (2014-2017). This strong relationship is further supported by a relatively large first-stage F-statistic of 55.8.

Table 1.5: IV Estimates — ER Visits (Selected Conditions)

	Selected Conditions				Selected Placebos		
	Asthma	All Resp.	Hypertension	Heart Attack	All Conditions	Diabetes	Appendicitis
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Smoke (α)	0.046***	0.015**	0.018**	0.001	0.009	−0.005	0.008
	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)	(0.006)	(0.005)
Medicaid (β)	0.023**	0.030***	0.027***	0.021***	0.064***	0.028***	0.005
	(0.010)	(0.009)	(0.009)	(0.005)	(0.011)	(0.007)	(0.004)
Smoke × Medicaid (τ_{IV})	−0.013***	−0.008***	−0.008***	−0.0002	−0.004**	0.001	−0.002
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Observations	543,052	543,052	543,052	543,052	543,052	543,052	543,052
Adjusted R^2	0.771	0.928	0.555	0.382	0.977	0.680	0.335

Notes: *p<0.1; **p<0.05; ***p<0.01. The dependent variable is the log of the count of ER admissions for the stated condition in ZIP j and week t of year y. Medicaid-related estimates have been scaled to reflect a ten percentage point increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across ZIP-weeks (equal to 17.2% coverage) or greater. All models include all of the fixed effects and controls listed in Table 1.4. Standard errors are clustered by ZIP.

In Table 1.5, I report diagnosis-specific estimates based on my preferred IV specification (column (6) of Table 1.4), with the asthma results reproduced in column (1) for comparison to the other health conditions. Columns (2)-(4) report results for several conditions that are well known to be exacerbated by air pollution, while the conditions in columns (5)-(7) are placebo conditions that in general should not be affected by poor air quailty. The effect of smoke and health insurance on ER visits for all respiratory-related conditions (column (2)) and hypertension (column (3)) are similar in magnitude to the asthma estimates. Comparing asthma visits to all respiratory visits, of which asthma is a subset, it is not surprising that the smoke effect (α) and the interaction effect (τ_{IV}) are both stronger when restricting specifically to asthma-related visits. The all respiratory category contains other respiratory conditions, such as influenza and pneumonia, which are more likely to be seasonal and to have fewer targeted preventive care options than are available to treat asthma. The estimates for heart attacks are of similar direction as the asthma estimates, though of smaller magnitude and not statistically significant. This may reflect the fact that ER visits for heart attacks are much less common than for asthma attacks, particularly when disaggregated to the ZIP-by-week level. Alternatively, it may reflect the possibility that, when triggered by wildfire smoke specifically, heart attacks more often result in a death than an ER visit. This possibility is supported by strong and statistically significant effect estimates related to heart attack mortality (as opposed to morbidity), which are presented in the following section. Each of the four selected pollution-related conditions has fairly similar estimates for the price and smoke effects, which in general are offset by an interaction (i.e., prevention) effect that is 30-50% in relative magnitude. Taken together, these results indicate that health insurance can have a strongly mitigating influence on pollution-related morbidity, without incurring a huge cost burden due to the increase in overall utilization when prices fall.

Results for the placebo conditions Table 1.5 help to validate the model and specifications, in addition to the use of smoke as a reduced form measure of air pollution. As we would expect, none of the placebo conditions has a statistically significant relationship with smoke exposure. In column (5) (All Conditions), the β estimate of 6.4% reflects the overall increase in the ER visit rate (for any medical condition) that we would expect to be caused by a ten percentage point increase in Medicaid coverage. The associated prevention effect of -0.4%

is significantly smaller in magnitude, most likely since the vast majority of acute health conditions requiring emergency medical attention are not related to smoke. Notably, there is no price effect observed for appendicitis – this condition is generally understood to strike individuals at random, and must be treated immediately due to the excrutiating pain it often typically (as well as the high risk of death if left untreated). For these reasons, we would not expect ER demand for appendicitis to be highly sensitive to price. Overall, the placebo models confirm that the smoke-related estimates obtained for the selected conditions of interest are indeed likely to have been triggered specifically by poor air quality.

1.6.2 Mortality

In most analyses of the effect of air pollution on health, mortality accounts for the largest share of overall costs. In this section, I consider the effect of health insurance and smoke on diagnosis-specific mortality, reporting IV estimates for pollution-related and placebo conditions in Table 1.6. With mortality, the main Medicaid effect can no longer be interpreted as a price effect, as we would not expect mortality to rise when the price of health care falls, as discussed in Section 1.3. In the mortality context, β can be interpreted as an additional prevention effect which operates only when the smoke indicator equals zero. Similar to the morbidity analysis, to simplify the interpretation of the effect estimates, $Smoke_{jty}$ is constructed to be a binary variable equal to one standard deviation of the smoke distribution or greater. For the county-by-month mortality analysis, this is 15.4% coverage.[23]

Few of the mortality estimates in Table 1.6 are statistically significant, which is not unexpected given that mortality events are relatively rare occurrences as compared to ER visits. Also, due to data constraints, the mortality analysis is conducted using far fewer records than the morbidity analysis, by a factor of over one thousand (i.e., due to the county-by-month aggregation as opposed to ZIP-by-week). However, heart attack-related mortality, reported in column (4), shows a strong response to smoke, both in its association with smoke alone and with the smoke and Medicaid interaction. Similar in relative magnitude to the analysis of asthma-related ER visits, the preventive effect of Medicaid reduces the mortality burden of smoke by over 30%, in this case with no corresponding price effect. The large magnitude of this result for heart attacks is somewhat surprising, but likely refects the low baseline rate of heart attack deaths in a given county-month.[24] As hypothesized, the main Medicaid effect is not statistically significant for any of the reported conditions. The interaction effect τ_{IV} is negative for nearly all of the conditions, although with limited statistical precision for all of the causes of death other than heart attacks. Overall, these estimates suggest a potentially important role for health insurance in the prevention of wildfire smoke-related mortality, but standard errors are wide enough that I cannot reject the null of no effect at the 90% confidence level, except in the case of heart attacks.

[23] I conduct a supplemental analysis of ER visits at the county-by-month level for comparison purposes, and the results are not substantially different from the original morbidity analysis at the ZIP-by-week level. Results from this supplemental analysis are contained in Appendix Table A.11.

[24] Across all counties and years, the median number of heart attack-related deaths in a given county and month is 2.2 per 100,000 residents. For comparison, the rate for respiratory-related deaths is 5.4 per 100,000 residents.

Table 1.6: IV Estimates — Mortality (Selected Conditions)

	Selected Conditions				Selected Placebos		
	Asthma	All Resp.	Hypertension	Heart Attack	All Conditions	Diabetes	Appendicitis
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Smoke (α)	0.106	0.115	−0.012	0.143*	0.041	0.041	0.027
	(0.077)	(0.086)	(0.093)	(0.080)	(0.062)	(0.078)	(0.028)
Medicaid (β)	0.002	0.034	−0.023	0.116	0.042	−0.091	0.002
	(0.059)	(0.058)	(0.112)	(0.120)	(0.055)	(0.141)	(0.022)
Smoke × Medicaid (τ_{IV})	−0.023	−0.034	0.006	−0.047*	−0.004	−0.002	−0.005
	(0.021)	(0.026)	(0.028)	(0.025)	(0.019)	(0.023)	(0.008)
Observations	4,176	4,176	4,176	4,176	4,176	4,176	4,176
Adjusted R^2	0.925	0.950	0.916	0.916	0.989	0.900	0.142

Notes: *p<0.1; **p<0.05; ***p<0.01. The dependent variable is the log of the count of deaths for the stated condition in county j and month t of year y. Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across county-months (equal to 15.4% coverage) or greater. All models include county-month-of-year and year fixed effects and weather and socioeconomic controls. Standard errors are clustered by county.

1.6.3 Treatment Effect Heterogeneity

The preceding sections measure the effects of health insurance and smoke on health outcomes for the complier population as a whole. However, previous research on air pollution and health has shown that impacts of exposure can be heterogeneous across observable and unobservable factors, including individual characteristics, regional socioeconomic trends, and background exposure levels. The mortality data collected by CDC contains a rich set of individual-level characteristics which I rely on in this section to investigate heterogeneity in mortality impacts across age group and broad race category. Table 1.7 presents these results for asthma-related mortality. Each column represents a separate regression in which the outcome variable is the log of the count of asthma deaths in a county and month for the relevant age or race group. Although the age-based estimates are not statistically significant, the main (and interaction) effect of smoke increases monotonically with age. This pattern is consistent with previous research: less healthy individuals tend to be more sensitive to smoke, and age is often a reasonable proxy for health.[25] Notably, the race-specific regressions for asthma-related mortality in Table 1.7 suggest that Blacks are both significantly more harmed by smoke exposure, and are also the most benefited by a public insurance expansion. The estimates imply that the asthma-related mortality rate for Blacks in a given county and month increases by 11.5% when smoke coverage is at least 15.4% or higher during the month, but that this health burden is reduced by 3.3% for a ten percentage point increase in the Medicaid penetration rate. Similar to the overall morbidity and mortality results, this represents a health insurance-induced reduction in the overall burden of smoke exposure of about 30%.

[25]California's Medicaid expansion primarily improved health care access for ages 19-64, but seniors are eligible for supplemental Medicaid insurance on top of Medicare benefits if they meet the same income requirements.

Table 1.7: IV Estimates — Mortality (Asthma), by Age and Race Groups

	Age Group			Race Group			
	1-18	19-64	65+	White	Black	Hispanic	Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Smoke (α)	0.012	0.039	0.102	0.053	0.115**	0.061	0.004
	(0.017)	(0.085)	(0.078)	(0.084)	(0.050)	(0.069)	(0.089)
Medicaid (β)	0.016	0.064	0.002	0.010	0.009	0.015**	0.009
	(0.026)	(0.063)	(0.056)	(0.066)	(0.042)	(0.006)	(0.053)
Smoke × Medicaid (τ_{IV})	−0.004	−0.005	−0.019	−0.006	−0.033**	−0.012	−0.006
	(0.005)	(0.026)	(0.022)	(0.023)	(0.015)	(0.021)	(0.025)
Observations	4,176	4,176	4,176	4,176	4,176	4,176	4,176
Adjusted R^2	0.132	0.762	0.919	0.911	0.836	0.801	0.804

Note: *p<0.1; **p<0.05; ***p<0.01. Each column represents a separate regression. *p*-values have been adjusted for multiple testing using the Bonferroni correction. The dependent variable is the log of the count of asthma deaths for the given age or race group in county j and month t of year y. Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across county-months (equal to 15.4% coverage) or greater. All models include county-month-of-year and year fixed effects and weather and socioeconomic controls. Standard errors are clustered by county.

Results for heart attack-related mortality (reported in Table 1.8) are similar to those for asthma, but for this disease the largest effects manifest among Hispanics. On average, Hispanics are much more affected by smoke than any other race, and a Medicaid expansion of ten percentage points reduces the mortality burden by just over 30% when smoke coverage crosses the one standard deviation threshold. As with asthma mortality, the main and interaction effects of smoke monotonically increase with age, although the estimates are imprecisely estimated due to the relatively small sample sizes. It is unfortunately not possible using the current research design to investigate the underlying causal mechanisms for these race-based results, but examination of relevant characteristics (e.g., race and income) across levels of the policy-exposure instrument (as discussed in Section 1.2.2) indicates that there exists sufficient variation in exposure to identify differences in impacts across these demographic groups. For example, if there was minimal representation of a certain racial group in the highest exposure areas then these race group-specific estimates might lack sufficient statistical power to be comparable across groups, but the statistics in Table 1.2 indicate that lack of coverage is likely not a significant concern. Moreover, the *California Health Interview Survey* of new Medicaid enrollees in 2015 found no statistically significant differences in program take up across race groups, which is further reassuring. Without leaning too heavily on these race-based results, given statistical imprecision and unclear mechanisms, these estimates are at least weakly suggestive that public health insurance might be a useful policy lever to counteract the negative health impact of poor air quality for two historically marginalized populations.

Table 1.8: IV Estimates — Mortality (Heart Attacks), by Age and Race Groups

	Age Group			Race Group			
	1-18	19-64	65+	White	Black	Hispanic	Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Smoke (α)	−0.003	0.066	0.118	0.098	−0.004	0.192*	−0.029
	(0.006)	(0.095)	(0.087)	(0.070)	(0.052)	(0.090)	(0.079)
Medicaid (β)	0.008	0.159*	0.068	0.133	0.026	0.059	0.081
	(0.008)	(0.079)	(0.122)	(0.117)	(0.052)	(0.097)	(0.063)
Smoke × Medicaid (τ_{IV})	0.001	−0.017	−0.042	−0.032	−0.004	−0.061**	0.003
	(0.002)	(0.030)	(0.028)	(0.021)	(0.016)	(0.028)	(0.023)
Observations	4,176	4,176	4,176	4,176	4,176	4,176	4,176
Adjusted R^2	0.040	0.789	0.908	0.892	0.822	0.831	0.849

Note: *p<0.1; **p<0.05; ***p<0.01. Each column represents a separate regression. *p*-values have been adjusted for multiple testing using the Bonferroni correction. The dependent variable is the log of the count of asthma deaths for the given age or race group in county j and month t of year y. Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across county-months (equal to 15.4% coverage) or greater. All models include county-month-of-year and year fixed effects and weather and socioeconomic controls. Standard errors are clustered by county.

1.6.4 Discussion

Relative estimates of the mitigating effect of health insurance on pollution-related health damages are fairly consistent across diagnoses, as well as across health outcomes (ER visits versus deaths). For diagnoses for which I estimate a statistically significant mitigating effect of health insurance on smoke-related ER visits and/or mortality – such as asthma, all-cause respiratory issues, hypertension, and heart attacks – the effect tends to be approximately 30-40% of the main smoke effect. Under the assumption that the Medicaid expansion only affected ER visits and mortality by expanding insurance coverage, then the estimated reductions in smoke-related ER usage and deaths can be directly interpreted as the treatment effect of a 10 percentage point increase in Medicaid coverage on each of these outcomes. However, as with any IV model, each estimate generated by this empirical model is a local average treatment effect (LATE) rather than a population-level average treatment effect. The LATE estimates from the IV-DID model capture the average effect of Medicaid on the subpopulation of Californians who gained coverage as a result of the Medicaid expansion *and* at the time of their health event (ER visit or death), live in a jurisdiction (a ZIP code or county) which had a high rate of uninsured adults in 2013. If the ACA was particularly effective at expanding Medicaid coverage among residents whose ER demand (or propensity to die) in response to smoke exposure is particularly sensitive to insurance (e.g., those with chronic and/or highly serious medical conditions), the estimated effects of Medicaid in this paper may be larger than corresponding average treatment effects for the entire uninsured population.

These estimates are broadly in line with previous work that has estimated the main (separate) effects of health insurance and air pollution on ER visits and mortality. For

example, Deschenes et al. (2017) estimate a 0.4% reduction in the all-cause mortality rate during months in which the *NOx Budget Program* was operating.[26] Although it is not statistically significant, this point estimate is essentially the same as the estimate I obtain for all-cause mortality for a ten percentage point Medicaid expansion, reported in column (5) of Table 1.6. It is also approximately the effect I estimate on heart attack-related mortality specifically, for a one percentage point increase in the Medicaid rate (column (4) of Table 1.6). The estimates in this paper are also consistent with previous work when accounting for the first-stage effect of wildfire smoke on the concentration of fine particulates, rather than simply considering the reduced form effect of the smoke. In versions of the main analyses in which I instrument for $PM_{2.5}$ concentrations with smoke coverage (Appendix Section A.6), the estimate of the first stage effect of smoke on mean weekly and monthly $PM_{2.5}$ concentrations is about 2.5 $\mu g/m^3$. Combining this with my estimate of a 0.4% increase in overall mortality when the air is smoky, this corresponds almost exactly to the Deryugina et al. 2019 estimate of a 0.2% increase in all-cause (elderly) mortality associated with each 1 $\mu g/m^3$ increase $PM_{2.5}$ concentrations. Overall, the percent increase in ER visits and mortality that I estimate for smoke is consistently in the low single digits, which is in the general range estimated in recent pollution-health studies that leverage other sources of quasi-random variation in pollution exposure.

Estimates of the price effect of insurance, β, are similarly in line with previous quasi-experimental work. For example, Card, Dobkin, and Maestas 2008 study the effect of Medicare – another public health insurance program – using a regression discontinuity design based on the strict age-65 eligibility threshold. These authors estimate a local average treatment effect of insurance for all-cause ER visits of 3.3%. Although they leverage a fundamentally different source of variation in public health insurance, and study a relatively older population, the estimate they obtain is similar in direction and magnitude to my estimate of 6.4%. A remaining open question is whether and to what degree the increase in ER utilization that accompanies an insurance expansion is inefficient.[27] Without additional information on patient-level characteristics and the underlying severity of illness that prompted the ER visit, it is difficult to ascertain whether these additional post-reform visits are socially optimal. However, comparing estimates of β for conditions with differing levels of presenting symptoms but a comparable risk of death without treatment, such as appendicitis and heart attacks, may shed some light. Appendicitis presents with increasingly severe and recognizable pain, whereas heart attack symptoms are often non-specific and may go ignored. The 0% price effect for appendicitis implies that the health insurance reform had no impact on ER demand for appendicitis. The 1.9% estimated health insurance impact for heart attack visits, given the severity of the condition and high likelihood of being a true emergency, suggests that these patients were under-utilizing emergency care relative to its efficient level in the pre-reform period. It is harder to make the same argument for the other pollution-related conditions, since these tend to have lower fatality risk than heart attacks. However, that the β estimate for each of these other pollution-related conditions is around 2% whereas the estimate for all conditions is over 6% implies that pollution-related

[26] One key difference between this paper and Deschenes et al. is that those authors link their estimated mortality impacts to reductions in ambient ozone, rather than to ambient $PM_{2.5}$.

[27] As discussed, this concern only applies to the ER visits analysis and not the mortality analysis, and it remains that case that mortality impacts by far comprise the largest source of estimated benefits.

health conditions are less prone to inefficient over-utilization of the ER than are conditions not affected by pollution. Taken together, these comparisons suggest that insurance-related moral hazard may be less of a concern in the context of using health insurance to adapt to recurring air quality shocks, as compared to standard health insurance models.

1.6.5 Counterfactual Analysis

In this section, I present monetized benefits from avoided ER visits and premature deaths based on IV-DID estimates of τ_{IV} under several counterfactual increases in the percentage point penetration of Medicaid coverage across regions within California. Specifically, I use the model to predict ER visits and mortality, using either the true Medicaid penetration rate in a given jurisdiction and year, or the counterfactual rate obtained after various percentage point increases in the true rate.[28] Predictions are made at the original level of analysis, ZIP-by-week (for ER visits) or county-by-month (for mortality), and then aggregated to produce a California-wide total health impact for each year, which are reported in Appendix Section A.7. It is important to note that the estimates of avoided health events are only for the ZIP-weeks and county-months where the smoke indicator $Smoke_{jty}$ is turned on, or where smoke coverage in jt is *at least* one standard deviation of the respective smoke distribution. Therefore, fluctuations in actual and predicted numbers of ER visits and mortality across years as reported in the table in part reflect true differences in the severity of wildfire seasons and associated smoke exposure across years.

Under the assumption of linearity, estimated impacts scale proportionally based on the counterfactual Medicaid increases, as is clear from the estimates in Appendix Table A.10. The assumption of linearity in Medicaid impacts is fairly strong, and comes with some important caveats. Medicaid enrollment may be correlated with underlying health status (or expected health risk), which would result in diminishing returns to insurance if there is negative selection in enrollment (e.g., if the sickest individuals are the first to enroll). However, Wang and Trivedi's analysis of the 2015 *California Health Interview Survey* noted that rates of asthma were not statistically different between enrollees and non-enrollees among the sample of newly eligible Californians. This provides some reason to rule out the potential for negative selection, at least if narrowly focused on air pollution-related health outcomes like asthma. Another potential threat to the linearity of impacts assumption is that there may be crowd out of other forms of insurance coverage as the Medicaid rate rises. While certainly possible, this is probably not highly representative of the Medicaid population in practice due to the strict eligibility threshold for annual income of 138% of the federal poverty line (approximately $16,000 for an individual). This income threshold is far below what would likely qualify for employer-sponsored insurance, and also likely too low for an individual to be able to comfortably afford unsubsidized monthly premiums for private insurance plans. Consequently, it seems quite likely that the majority of new Medicaid enrollees were previously uninsured, somewhat reducing concerns about crowd out. For these reasons, while linearity remains a strong assumption and the benefit estimates reported

[28] Counterfactual Medicaid rates are capped at one hundred percent in jurisdictions where the increase would result in more than full coverage (i.e., where the counterfactual number of Medicaid enrollees would be greater than total regional population).

in this section are meant primarily for illustrative purposes, typical concerns regarding a linearity assumption are less applicable to this setting.

I report benefit estimates for asthma-related ER visits and heart attack-related mortality only, since these health outcomes are where the protective influence of health insurance was most clear. However, this set of results represents only one small piece of a larger picture about how health insurance can reduce the health burden of air pollution exposure. Moreover, these estimates represent an even smaller piece of the overall benefits to society of public health insurance, the vast majority of which are almost surely unrelated to air pollution. The estimates in Panel (a) of Table A.10 indicate that a ten percentage point increase in Medicaid coverage results in approximately 500 fewer annual ER visits for asthma during years that are particularly smoky (e.g., 2016 and 2017). However, if accounting for the price effect β, we would see a net *increase* in ER visits for a given insurance expansion, causing the estimates to switch sign but be only about one third of their original magnitude based on the relative magnitudes of the coefficients in Table 1.5. The average cost of an asthma-related ER visit is about $1,500 (Wang et al. 2014), which implies an annual increase in the monetized public burden of ER visits for asthma of approximately $250,000 for a ten percentage point increase in Medicaid.

Estimated mortality impacts are not complicated by a price effect, and are significantly more meaningful in economic terms when monetized using standard values for lost life. Panel (b) of Table A.10 reports estimates of annual avoided heart attack deaths for a ten percentage point Medicaid expansion. Again focusing on the estimates in relatively smokier years, a ten percentage point increase in Medicaid in California reduces heart attack-related mortality by about 20-25 deaths per year. Applying the current average value of a statistical life (VSL) of $10,000,000, this implies mortality benefits – for heart attacks only – of 200 to 250 million dollars per year. Again, this represents a lower bound on overall insurance benefits, specifically related to air pollution, and further narrowed to focus on a single health condition. To put this estimate in context, EPA's 2012 Regulatory Impact Assessment of the national ambient $PM_{2.5}$ standard projected that about 90 premature deaths would be avoided annually by lowering the threshold for mean $PM_{2.5}$ concentrations by 3.0 $\mu g/m^3$ and assuming full compliance by all states (U.S. EPA 2011, p. ES-17). Given that EPA's estimate is for the entire country, as opposed to the California-only estimates in this paper, this suggests that the estimated reduction in pollution-related fatal heart attacks due to expanded public insurance in California is both of reasonable magnitude, and also large enough to be policy-relevant. Further, this comparison suggests that California's actual Medicaid expansion of about 10 percentage points delivered mortality benefits on the order of those projected to manifest under national environmental policy governing ambient air quality. To my knowledge, this is the first time that this comparison has been made.

1.7 Conclusion

Using California as a case study, this paper assesses whether access to public health insurance can mitigate the negative health impacts of exposure to air pollution. This is a policy-relevant question in the U.S. given recent and widespread public health insurance expansions enacted under the Affordable Care Act, and the likelihood of near-term air quality deterioration due

to both the dismantling of environmental regulations in recent years and increasingly severe wildfire seasons in the drought-stricken West. While the individual effects of air pollution and health insurance on health outcomes have been studied previously in great depth by researchers across disciplines, this paper is the first to consider the two effects in the same empirical framework.

My research strategy brings together quasi-random variation in air pollution exposure generated by wildfire smoke, and regional variation within California in exposure to the Medicaid program expansion that began in 2014. Smoke plumes from wildfires contain high concentrations of $PM_{2.5}$, a pollutant directly regulated by EPA and widely understood to be one of the most harmful to human health, and can travel unpredictably across wide distances and short time horizons based on wind and other weather patterns. Further, the impact of California's Medicaid expansion was experienced differently across regions based on the number of newly eligible individuals residing in each locality. I rely on this plausibly random spatial variation in exposure to the health care policy to address the classic problem of endogeneity in insurance take-up decisions using an instrumental variables research design.

My preferred estimates suggest that a one percentage point expansion in public insurance reduces the heath burden of smoke exposure by approximately 3%, a finding that is fairly consistent across smoke-related health conditions and health outcomes (i.e., ER utilization versus mortality). I estimate that a ten percentage point increase in Medicaid coverage – which approximates the actual post-ACA increase in California's Medicaid enrollment, as well as the percentage of uninsured remaining in the state as of 2020 – has resulted in around 20 fewer heart attack deaths and 1,500 asthma ER visits annually. Together, using standard values for lives lost and costs of ER hospitalization, these avoided health impacts are valued in the hundreds of millions of dollars each year, and yet represent only a small fraction of public insurance benefits related to two (of potentially many) health outcomes that might affected by smoke. Moreover, these effects represent reductions in mortality and hospitalizations over and above the health benefits of public insurance when air pollution is minimal.

A common concern of policymakers regarding public health insurance is increased utilization of health care services following a decrease in price, a phenomenon which I indeed document in this setting. The concern centers around whether or not the post-reform increase in utilization is socially "efficient" – in that it represents truly necessary care. By comparing estimates of this effect across different health conditions in the California Medicaid setting, I provide suggestive evidence that the observed increases might be socially efficient. However, more work is needed to examine this issue further, ultimately to determine whether the benefits of public health insurance, including the novel air pollution-related benefits that are estimated in this paper for the first time, outweigh overall costs.

Wildfires and associated smoke plumes pose a continued threat to population health in western regions of the U.S. that have seen persistent drought in recent years. California wildfire seasons in 2017 and 2018, and again in 2020, have broken previous records for acres burned, and sent plumes of smoke and ash hundreds of miles away from initial ignition sites, thereby exposing much of the state's population to harmful particulate pollution over sustained periods. By demonstrating the value of improved access to health care (through more widely available health insurance) in reducing the severity of health impacts from air pollution, this study's findings may have relevance for the mitigation of negative health

impacts linked to future wildfire events. Further, given the strong association of wildfire smoke with $PM_{2.5}$, a harmful pollutant that has been regulated by EPA for decades, the findings of this paper may have implications for addressing the persistent health burden of air pollution more generally. Despite this clear takeaway, this paper has some important limitations. Specifically, in addition to further exploring the potential issue of over-utilization of the ER, understanding the causal mechanisms for the effects that I have documented – whether the benefits come through improved take up of preventive care, or something else not yet identified – is a clear area for future research.

Chapter 2

Declining Power-Plant Emissions, Co-Benefits, and Regulatory Rebound

2.1 Introduction

Over the past half-century, expansive regulations have been introduced under the Clean Air Act to reduce emissions of harmful air pollution. Hundreds of studies have assessed the benefits of these regulations, the vast majority of which manifest as reductions in premature mortality. In recent years, 'co-benefits' have accounted for a substantial share of the assessed benefits from EPA emissions regulations. This increased prominence has invited increased scrutiny of the indirect impacts of emissions regulations. This paper investigates co-benefits accounting in theory and practice, bringing new evidence to bear on complex causal relationships between reductions in power plant emissions, improvements in downwind air quality, and associated health outcomes.

An environmental regulation generates co-benefits when it delivers not only the environmental improvements targeted by the authorizing legislation, but also reductions in other related damages. For example, regulations targeting greenhouse gas emissions (GHGs) may induce investments that reduce not only GHGs, but also precursors to the formation of small particulate matter ($PM_{2.5}$). Figure 2.1 shows how co-benefits from reduced exposure to $PM_{2.5}$ have comprised a significant share of assessed benefits under recently proposed climate-change-mitigation efforts (e.g. the Clean Power Plan and ACES) and other electricity sector regulations (e.g. the Mercury Rule). Looking ahead, the local pollution impacts of the accelerated deployment of low-carbon resources will have a substantial role to play in policy discourse surrounding climate change.

Accounting for both the direct and indirect effects of a policy or market intervention is consistent with fundamental economic principles of benefit cost analysis (see, for example, Gramlich, 1990; Boardman et al. 2018). Guidance issued by the Office of Management and Budget (OMB) is unequivocal, stating that agencies conducting benefit cost analysis should "[i]dentify the expected undesirable side-effects and ancillary benefits of the proposed regulatory action and the alternatives. These should be added to the direct benefits and costs as appropriate" (OMB 2003, pp. 2-3). What is less clear is *how* to properly account for indirect versus direct impacts.

Figure 2.1: Co-benefits as a Share of Total Projected Benefits

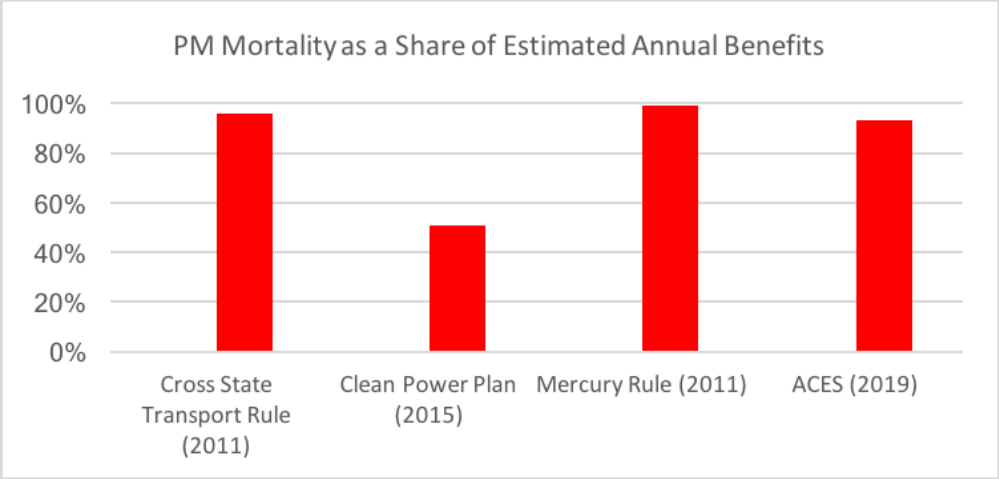

Notes: Estimated "co-benefit" from reduced particulate matter ($PM_{2.5}$) exposure account for over 90% of mortality benefits from recent regulations targeting greenhouse gas (GHG) emissions, mercury emissions. Source: EPA Regulatory Impact Assessments.

In cases where markets operate efficiently, indirect impacts in 'secondary' markets can be safely ignored when they are captured by measured surplus changes in primary markets (Boardman et al. 2018). Additional measurement and accounting of indirect impacts becomes more important in the presence of market failures such as environmental externalities. The increased emphasis on health co-benefits in the assessment of emissions regulations has prompted increased scrutiny around the accounting of these indirect benefits. The EPA Science Advisory Board recently noted the absence of "any careful scientific guidance on how co-benefits should be identified and estimated in regulatory analyses" (U.S. EPA, May 2019).

This paper advances the theory and informs the practice of co-benefits accounting in the context of air pollution regulation. We formalize a framework for analyzing indirect benefits in the presence of overlapping air quality regulations. Using two recent emissions policy regimes as case studies, we implement this framework empirically. We leverage high-resolution meteorological data, recent advances in the modeling of air particle trajectories, and significant variation in power plant emissions across the United States over the period 2005-2019. We integrate simulated air particle trajectories into an empirical analysis of how reductions in coal plant emissions impact downwind $PM_{2.5}$ concentrations.

Empirical estimates of how changes in power plant emissions have impacted ex post observed $PM_{2.5}$ concentrations are useful for evaluating the economics of co-benefits along a number of dimensions. First, we systematically compare our empirical estimates with the projections of workhorse models that are widely used in ex ante and ex post assessments of pollution co-benefits. Second, we investigate the extent to which policy interactions generate discrepancies between projected and realized co-benefits. In the assessment of pollution co-benefits, it is standard to assume away any meaningful interactions between new and existing policies. We note the potential for 'regulatory rebound' in areas where upwind reductions in power plant emissions relieve the constraints that were previously imposed by threshold-based air quality standards. We propose an empirical test for this local regulatory response

which effectively trades off local air quality improvements and locally incurred pollution control costs. Finally, to the extent that we document significant discrepancies between standard accounting predictions and realized air quality impacts, we incorporate ex post observed mortality outcomes to assess the health implications of these discrepancies.

In this working paper, we report on an preliminary exercise that uses a less computationally intensive approach to modeling pollution fate and transport. We systematically compare the results of these simpler simulations against $PM_{2.5}$ monitor readings in downwind locations. This exercise provides some initial insights into the accuracy with which standard models predict observed air quality outcomes. It also supports a very preliminary empirical test for regulatory rebound in areas constrained by local air quality standards. We find weak but suggestive evidence of rebound in local pollution levels following sustained emissions reductions that were prompted by several recent emissions regulations. Based on our current transport modeling and empirical strategy, estimates of emissions reductions on long-run mortality are also weakly suggestive of rebound. In future work, we will incorporate more precise particle trajectory modeling which is better suited to this setting.

The paper proceeds as follows. In Section 2.2, we review the historic decline in power plant emissions that has occurred over the past twenty years and describe the implications of several recent regulatory interventions that we use as case studies in our analysis. In Section 2.3, we summarize the role of $PM_{2.5}$ in the quantification of co-benefits under emissions policy proposals and we outline our theory of regulatory rebound with respect to existing air quality standards. In Section 2.4 we describe our current method for modeling the transport of emissions particles from power plants to downwind locations and the formation of precursors into $PM_{2.5}$, and Section 2.5 describes our data sources. Section 2.6 presents and discusses our empirical results and Section 2.7 concludes.

2.2 Power Plant Emissions

Figure 2.2 documents a dramatic decline in air pollution emissions from US coal-fired power plants. The top panel plots sulfur dioxide (SO_2) emissions. The bottom panel shows nitrogen oxide (NO_x) emissions. There are a number of factors that could explain this downward trend, including increased competition from domestic shale gas, increases in renewable electricity generation, lower than expected energy demand, and EPA regulations that directly target emissions from power plants.

Figure 2.2: SO_2 and NO_x emissions, 2005-2019

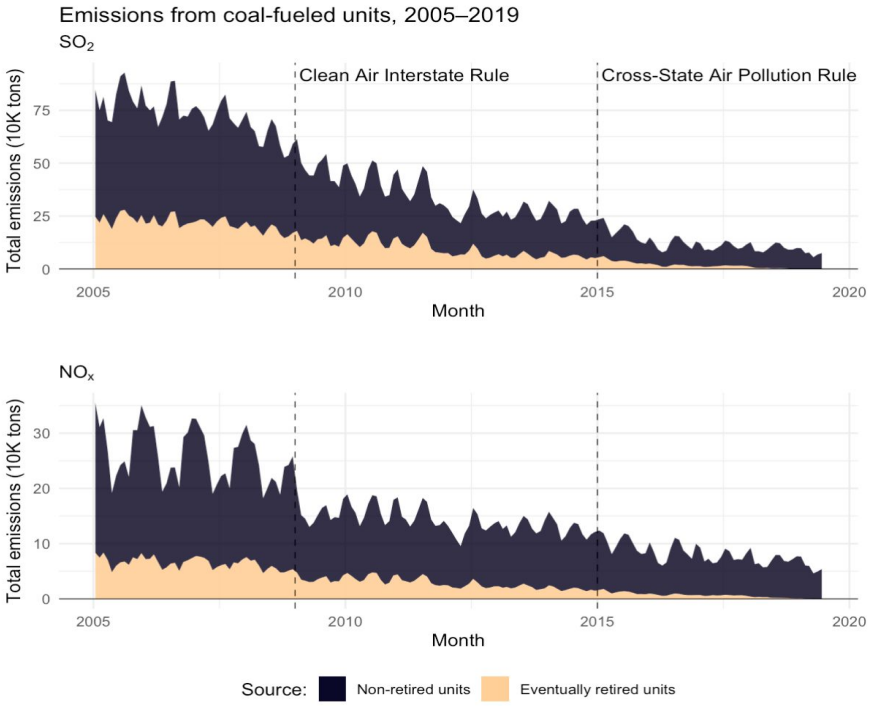

Figure 2.2 tracks coal plant emissions separately for plants that retired over this period and those that operated over this duration. The figure shows that plant retirements over this period explain a relatively small share of total reductions. A recent paper decomposes the decline in power plant emissions damages over the period 2010-2017 in terms of the scale of electricity production; the composition of power plants generating electricity; and the technology or emissions intensity of plants in the market (Holland et al. 2020). These authors find that the most significant driver of pollution damage reductions over this period is plant-level 'technique' changes that have reduced the emissions intensity of coal-fired generation. Technique changes have primarily involved capital-intensive pollution control technologies. Fuel switching has played a smaller role.

Regulations limiting power plant emissions have been the driving force behind the technique changes that have delivered significant reductions in the emissions intensity of coal-fired electricity generation. Over the fifteen year period we consider, the stringency of regulations limiting SO_2 and NO_x emissions increased substantially. The vertical lines in Figure 2.2 denote two distinct regulatory regimes. The first line denotes the implementation of the Clean Air Interstate Rule (CAIR) which expanded the scope of NO_x regulations and increased the stringency of SO_2 emissions regulations. CAIR introduced three different cap-and-trade programs: an annual program for SO_2, an annual program for NO_x and a summertime program for NO_x. Although this rule was announced in 2005, programs did not start operating until 2009.[1]

[1]A 2008 court decision upheld CAIR but instructed the EPA to replace the rule with another, which eventually became the Cross State Air Pollution Rule (U.S. EPA, 2015c, 2008c). The court remanded CAIR "without vacatur" because it found that allowing CAIR to remain in effect until replaced would temporarily preserve environmental values.

The Cross-State Air Pollution Rule (CSAPR) replaced CAIR in 2015 and it remains in effect today. This regulation imposes annual SO_2, annual NO_x, and ozone-season NO_x emissions "budgets" at the state level. 2015 was also the year that coal plants started complying with the Mercury and Toxics Rule (MATS). Because technologies that reduce mercury emissions also reduce other criteria air pollution, it was projected that MATS would reduce SO_2 emissions by more than 40% beyond the reductions expected from the Cross State Air Pollution Rule.

This sequence of increasingly stringent EPA regulations has delivered significant and sustained reductions in NO_x and SO_2 emissions from coal plants. Ex ante and ex post analyses associate large health benefits with these emissions reductions. Ex ante regulatory impact assessments of CSAPR and MATS projected annual benefits in the range of \$120-\$280 billion and \$25-\$70 billion, respectively. In an ex post analysis of emissions reductions over the period 2010-2017, Holland et al. 2020 estimate that coal plant investments in pollution abatement and fuel switching delivered annual health benefits in excess of \$60 billion.

The health benefits assessed in both ex ante and ex post studies are comprised almost entirely of health benefits from reduced exposure to small particulate matter ($PM_{2.5}$). The causal link between reductions in power plant emissions and reduced $PM_{2.5}$ exposure is complex and indirect. Complicated interactions between atmospheric chemistry, meteorology, topography, and pre-existing levels of precursors determine how changes in emissions of NO_x and SO_2 at an electricity generating unit in one location cause changes in ambient $PM_{2.5}$ concentrations at other locations. Ex ante projected health benefits, and many ex post benefits assessments, rely on complex integrated assessment modeling to map coal plant emissions changes to local air quality changes to health improvements. The validity of these projections rely to a significant extent on the accuracy with which these complex relationships are represented.

The complex modeling that underpins the assessment of co-benefits has come under some scrutiny (Aldy et al., 2020). In what follows, we leverage significant and staggered reductions in coal plant emissions, advances in high-resolution particle transport modeling, and daily monitoring of local $PM_{2.5}$ concentrations across thousands of locations to empirically investigate the relationships between coal plant emissions reductions and downwind air quality.

2.3 Co-benefits, Overlapping Policies, and Regulatory Rebound

In a recent review of benefit-cost analyses (BCA) of recent EPA air-pollution regulations, Aldy et al. (2020) highlight a number of striking facts. First, reduced exposure to small particulates ($PM_{2.5}$) has been the most economically significant source of prospective benefits from major air-pollution regulations. Second, "co-benefits" (i.e., benefits caused—but not directly targeted—by CAA regulations) comprise almost half of the monetized benefits assessed in prospective regulatory assessments. Third, reductions in exposure to $PM_{2.5}$ represent 96 percent of all monetized co-benefits over this period. For example, $PM_{2.5}$ co-benefits account for more than 99 percent of benefits assessed under MATS.

As we note in the prior section, complex relationships between air pollution emissions,

$PM_{2.5}$ formation, and dose-response relationships complicate the assessment of local air pollution co-benefits. This assessment exercise is further complicated by the potential for interactions between overlapping regulations. We consider, in particular, interactions between regulations limiting air pollution from coal plants (e.g. CAIR, CSAPR, MATS) and the ambient air quality standards that specify maximum allowable air pollution concentrations for air pollutants such as $PM_{2.5}$ and ozone. These health-based national ambient air quality standards (NAAQS) apply uniformly across the country, and have been a central regulatory mechanism governing air quality improvements across the United States in recent decades.

Threshold-based NAAQS are enforced using pollution concentration measurements collected across the EPA's air pollution monitoring network. Compliance with these national standards is assessed at a local level on the basis of 'design values' (DVs) which are calibrated using local air quality monitor readings. Within three years of a new or revised NAAQS, states must submit state implementation plans to demonstrate how they will comply. The SIP process relies on emissions inventories and air quality modeling to determine the portfolio of control strategies (such as mandating abatement investments or limiting entry of new sources) required to achieve compliance. If a region fails to comply with NAAQS, it must take additional steps to improve air quality.

It is important to note that emissions from distant upwind sources can significantly affect a county's design value (and thus NAAQS compliance status). Recent research has demonstrated that a majority of local PM levels measured at EPA monitors had originated at sources outside the home county (see, for example, Wang et al. 2020). This and related findings underscore that local regulators only control a fraction of emissions that contribute to local air quality. A review of approved SIPs finds that states routinely account for national or regional emissions regulations in their compliance strategy planning. In contrast, prospective regulatory impact analysis does not.

Figure 2.3 provides a stylized illustration of pollution damages and abatement costs in a jurisdiction that is constrained by an ambient air quality standard. The vertical black line represents the pollution concentration threshold that localities should not exceed under NAAQS. In the top panel, local pollution levels exactly meet this threshold. The downward sloping line represents the local marginal abatement cost curve. The area A represents the costs of remaining in compliance with the standard.

The horizontal line represents marginal damages caused by increased pollution exposure. Consistent with the scientific assessment of health impacts from particulate matter exposure, we assume the damage function is locally linear. Figure 2.3 depicts a case where levels of abatement under the standard are inefficiently low given the pollution damage costs. A single national ambient air quality standard applies to all counties across the United States. Even if these standards are set efficiently on average, there will be many localities in which marginal abatement costs do not equal marginal damages at the threshold.

Consider an emissions regulation that would reduce emissions from large point sources which contribute to the formation of pollution summarized in Figure 2.3. Suppose this regulation would induce reductions in upwind sources such that local pollution levels fall to point G. Now that air quality standards are no longer binding, a rational response on the part of the local regulator in this location would be to release constraints on local emissions sources. Shapiro and Walker 2020 study local markets for pollution "offset" permits that are created when a new entrant must pay incumbent plants to reduce emissions and indeed

Figure 2.3: Policy Interactions and Regulatory Rebound

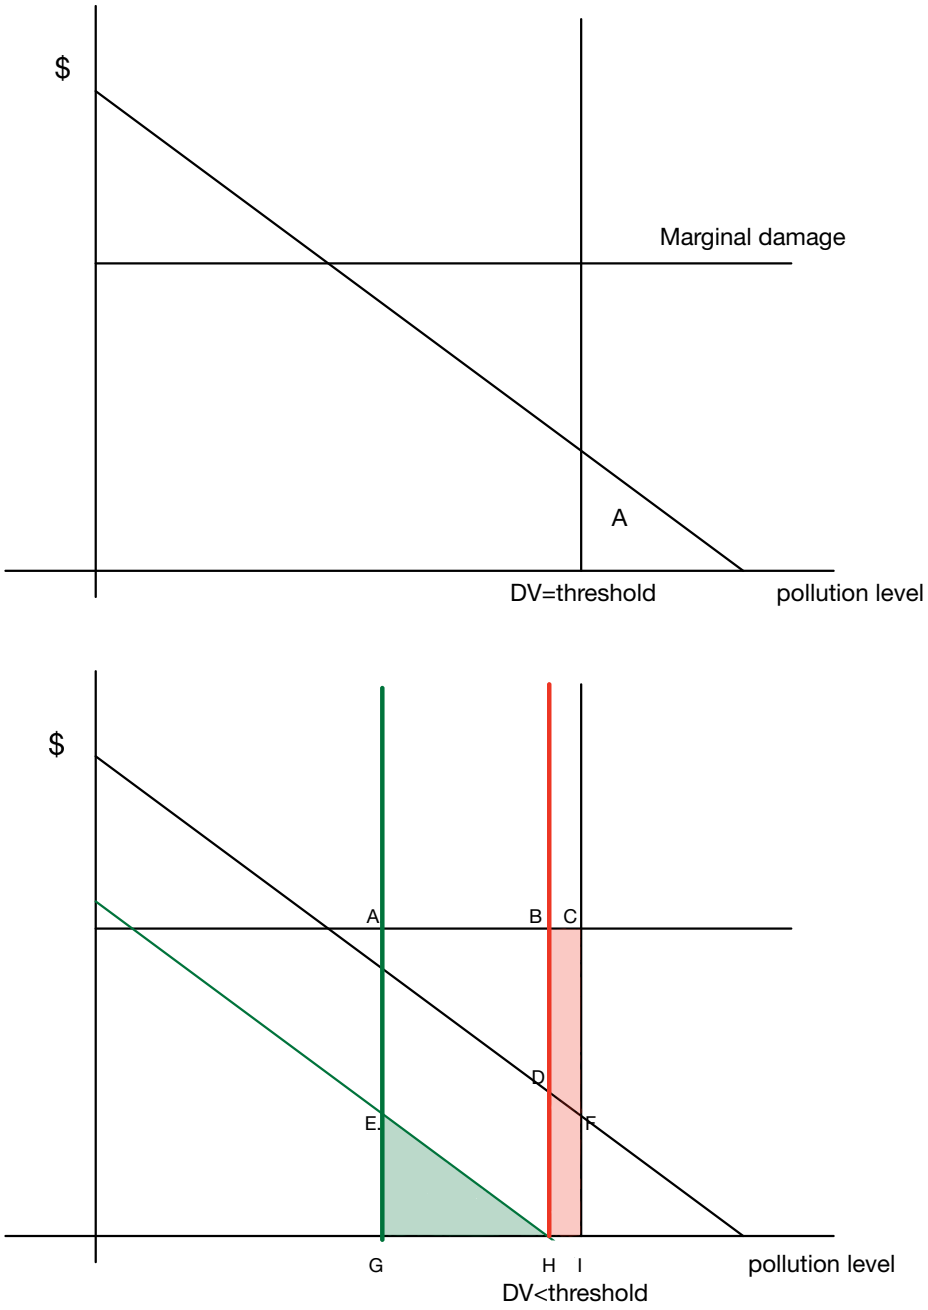

Notes: This figure provides a stylized illustration of 'regulatory rebound'. The top panel illustrates the abatement costs incurred locally to comply with a binding air quality standard. In the bottom panel, a reduction in regional emissions reduces local pollution concentrations. The local regulatory response is to permit local emissions up to the level that corresponds to zero abatement costs. This effectively foregoes health benefits (equal to area ABGH) in exchange for a reduction in local abatement costs of EGH.

show that the shadow value of air quality standards can be quite high in constrained areas. This 'regulatory rebound' in response to the new emissions regulation would increase local pollution levels to point H. In other words, increases in local pollution would counteract regional reductions induced by the new regulation.

We are not the first to highlight the potential for interactions between new and existing policies. For example, in its assessment of the MATS rule, the EPA notes that the proposed rule could "lead to reductions in ambient $PM_{2.5}$ below the NAAQS for PM in some areas and assist other areas with attaining the PM NAAQS" (U.S. EPA 2011a, pp. 5-2). But standard co-benefits accounting ignores the economic implications of this regulatory response. In Figure 2.3, standard approaches to assessing co-benefits ex ante would would project benefits represented by the area ACGI. After accounting for the local regulatory response, *realized* health co-benefits would equal area BCHI. Additional benefits accrue as increased economic activity (area EGH).

In this illustrated example, standard accounting practices over-estimate realized realized co-benefit value. But this need not be true in general. In constrained areas where marginal abatement costs exceed marginal damages from pollution exposure, failing to account for regulatory rebound could under-estimate realized co-benefits. The net implications will depend on how the economic value associated with increased local pollution compares to the foregone health co-benefits.

Despite the pivotal importance of projected $PM_{2.5}$-related co-benefits to the benefit-cost assessment of proposed emissions regulations, we are unaware of any ex post efforts to empirically quantify realized benefits of previous regulations. In a similar vein, standard integrated assessment modeling (*e.g.*, the InMAP model, which we rely on in this paper to simulate the transport of emissions particles from sources to downwind receptors) does not account for potential policy interactions or other factors which could result in actual pollution concentrations diverging from predicted levels, such as the regulatory-rebound phenomenon we describe. In what follows, we will estimate the impacts of recent emissions regulations on ambient $PM_{2.5}$ levels across the EPA monitoring network, as well as changes in long-run mortality rates. In an empirical test for regulatory rebound, we will also assess the extent to which these impacts appear to vary across regions that are more or less constrained by pre-existing ambient air quality standards (i.e., NAAQS).

2.4 Modeling PM Formation

In this working paper, we report results based upon the Intervention Model for Air Pollution (InMAP) from Tessum, Hill, and Marshall 2017, which is one of the reduced-form particle transport models that simulates both the transport of and transformation of precursor particles into ambient $PM_{2.5}$.[2] The InMAP project provides a reduced complexity approach that maps emissions to $PM_{2.5}$ concentrations using a calibrated source-receptor matrix approximation. More precisely, InMAP uses pre-processed physical and chemical information from detailed chemical transport modeling to calibrate a source-receptor matrix that can

[2]In subsequent versions of this paper, we will report on results based on particle transport simulations instead using the Hybrid Single-Particle Lagrangian Integrated Trajectory (HYSPLIT) model from the National Oceanic and Atmospheric Administration (NOAA).

be used to quickly simulate how emissions at one location impacts $PM_{2.5}$ concentrations at other locations (Tessum et al., 2017). InMAP takes as inputs plants' geographic coordinates, stack heights and diameters, and levels of precursor emissions—both NO_x and SO_2.

While InMAP includes quarter of year (or annual) source-receptor matrices, we run InMAP for each month of our sample—combining monthly total emissions with the month's respective quarter. Importantly, InMAP's quarterly source-receptor matrices are fixed across all years, thus abstracting away from potentially significant meteorological variation over time that is not season-invariant. The output is a grid for which each "cell" includes a level of $PM_{2.5}$ in $\mu g/m^3$.[3] The two panels of Figure 2.4 illustrate examples of the output from InMAP for January and July in 2005. Specifically, Figure 2.4 depicts the modeled $PM_{2.5}$ concentrations based upon the aggregate emissions reported by all electricity generating units in the contiguous U.S. during the stated month.

Figure 2.4: InMAP outputted $PM_{2.5}$ concentrations

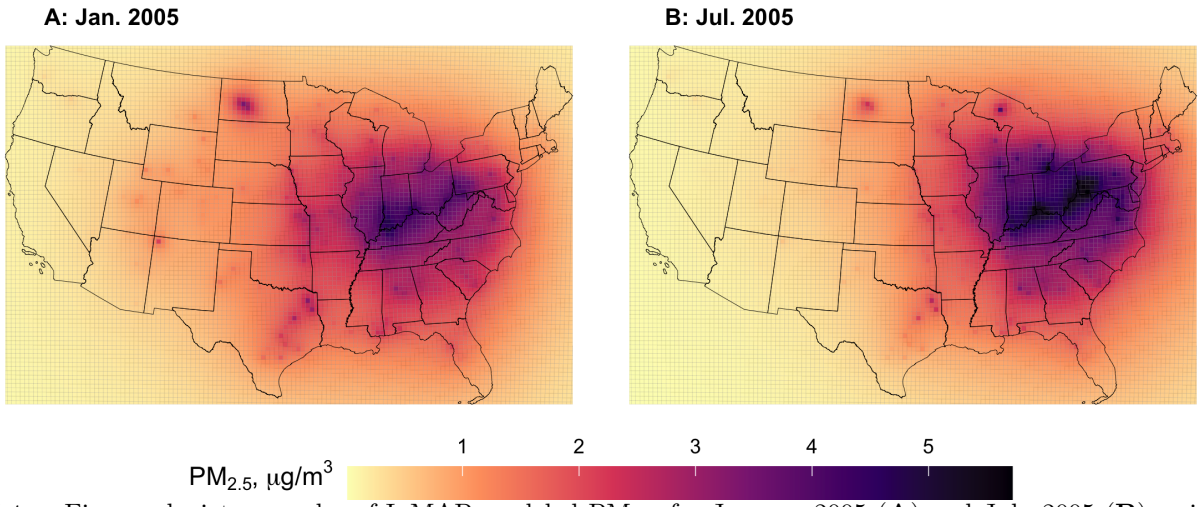

Notes: Figures depict examples of InMAP modeled $PM_{2.5}$ for January 2005 (**A**) and July 2005 (**B**), using aggregate emissions reported by all generating units in the CEMS database in the respective month.

Our first empirical exercise is to assess the extent to which InMAP predictions correspond to actual $PM_{2.5}$ measurements over our time period (2005-2019). Equation (2.1) summarizes our basic empirical approach for modeling this relationship:

$$(\text{Monitor PM}_{2.5})_{i,t} = \beta \left(\widehat{\text{PM}_{2.5}}\right)_{i,t} + \theta_i + \eta_{r(i),t} + \varepsilon_{i,t} \tag{2.1}$$

The dependent variable in equation (2.1) is the average concentration of ambient $PM_{2.5}$ measured at monitor i in time t. In this working paper, our analysis is conducted using data disaggregated to the monthly level, although in subsequent versions of this project we may shift to a finer temporal resolution, e.g. incorporating daily or hourly variation in emissions and pollution concentrations. $\widehat{\text{PM}_{2.5}}$ represents the InMAP-predicted level of ambient $PM_{2.5}$ at monitor i in month-of-sample t attributable to emissions of precursor pollutants by power plants in t. Importantly, these predicted concentrations represent $PM_{2.5}$ attributable to

[3]It is possible to get many other outputs from InMAP; in this paper, we are focusing on $PM_{2.5}$.

aggregate power plant emissions in these months, a large portion of which are driven by the plausibly exogenous policy changes described in Section 2.2. The β coefficient captures the statistical relationship between variation in these power plant emissions and local air quality, after accounting for monitor-specific effects (θ_i) and NERC region by month-of-sample trends ($\eta_{r(i),t}$). If InMAP is accurately modeling pollution transport and $PM_{2.5}$ formation, we should expect to estimate a β parameter equal to (or at least close to) one.

Next, we implement an empirical test for regulatory rebound. Intuitively, we want to assess whether localities that were constrained at baseline *take back* the gap between local $PM_{2.5}$ concentrations and air-quality standards following sustained, largely policy-induced drops in upwind emissions. If regulatory rebound does happen, then *upwind* decreases in $PM_{2.5}$ precursors, which are the primary inputs to our InMAP modeling, should reduce local $PM_{2.5}$ concentrations to a lesser extent when the locality was previously constrained by (or in the neighborhood of) an air-quality standard.

In order to test this hypothesis, we extend equation (2.1) by allowing the effect of *upwind* $PM_{2.5}$ to vary during the policy period by the monitor's baseline proximity to the NAAQS threshold for $PM_{2.5}$, i.e.,

$$
\begin{aligned}
(\text{Monitor } PM_{2.5})_{i,t} = \\
&\beta_1 \left(\widehat{PM_{2.5}}\right)_{i,t} \times \mathbb{1}\{\text{Far, Below}\}_{Baseline} + \\
&\beta_2 \left(\widehat{PM_{2.5}}\right)_{i,t} \times \mathbb{1}\{\text{Near, Below}\}_{Baseline} + \\
&\beta_3 \left(\widehat{PM_{2.5}}\right)_{i,t} \times \mathbb{1}\{\text{Near, Above}\}_{Baseline} + \\
&\beta_4 \left(\widehat{PM_{2.5}}\right)_{i,t} \times \mathbb{1}\{\text{Far, Above}\}_{Baseline} + \\
&\beta_5 \left(\widehat{PM_{2.5}}\right)_{i,t} \times \mathbb{1}\{\text{Far, Below}\}_{Baseline} \times \mathbb{1}\{\text{Policy Regime}\}_t + \\
&\beta_6 \left(\widehat{PM_{2.5}}\right)_{i,t} \times \mathbb{1}\{\text{Near, Below}\}_{Baseline} \times \mathbb{1}\{\text{Policy Regime}\}_t + \\
&\beta_7 \left(\widehat{PM_{2.5}}\right)_{i,t} \times \mathbb{1}\{\text{Near, Above}\}_{Baseline} \times \mathbb{1}\{\text{Policy Regime}\}_t + \\
&\beta_8 \left(\widehat{PM_{2.5}}\right)_{i,t} \times \mathbb{1}\{\text{Far, Above}\}_{Baseline} \times \mathbb{1}\{\text{Policy Regime}\}_t + \\
&\theta_i + \eta_{r(i),t} + u_{i,t}
\end{aligned}
\tag{2.2}
$$

where *Near* indicates that a monitor was within 1.0 $\mu g/m^3$ of the NAAQS for $PM_{2.5}$ in the baseline period of each policy regime, and *Far* indicates the opposite.[4] We currently focus on the annual NAAQS threshold for $PM_{2.5}$ but may also consider distance from the 24-hour $PM_{2.5}$ NAAQS in future extensions of this work.[5] Similarly, *Below* indicates that the monitor is below the standard, and *Above* indicates that the monitor's $PM_{2.5}$ level exceeds

[4]We define the baseline period for the CAIR regime (2009-2014) as 2007-2008 and the baseline period for the CSAPR/MATS regime (2015-2019) as 2013-2014. We average over design values in each baseline period to reduce noise from year-to-year pollution fluctuations.

[5]As an example of the NAAQS $PM_{2.5}$ limits: The annual standard was revised downward from 15.0 to 12.0 $\mu g/m^3$ (annual mean, averaged over three years) in 2012, with implementation beginning in 2014, and has stayed at that level through 2019, the last year of our panel.

the standard. Thus, $\mathbb{1}\{\text{Near, Below}\}_{Baseline}$ takes on a value of one if and only if monitor i is located in a region with a NAAQS design value that is between 0 and 1 $\mu g/m^3$ below the $PM_{2.5}$ standard in the baseline period. $\mathbb{1}\{\text{Policy Regime}\}_t$ takes a value of one in years when each policy is operating (2009-2014 and 2015-2019 for the CAIR and CSAPR/MATS regimes, respectively). The interactions between upwind PM and distance indicators in the pre-policy period (β_1 through β_4) account for pre-existing differences in particulate exposure from upwind sources that exist absent any regulatory action. In addition to monitor-level fixed effects, we also include an interaction for NERC region and month-of-sample. The latter two-way fixed effect is given by $\eta_{r(i),t}$ in equation (2.2). If regulatory rebound occurs in practice, then *upwind* decreases in $PM_{2.5}$ concentrations[6] during each emissions policy regime should reduce local $PM_{2.5}$ concentrations to a lesser extent when the locality is constrained (or in the neighborhood of) an air quality standard. In equation (2.2), we would expect to see no distance-based heterogeneity in the pre-policy period, so β_1 through β_4 should be of relatively equal size (and close to one, if InMAP is accurately predicting downwind $PM_{2.5}$). During each policy regime, regulatory rebound would manifest empirically as smaller pass-through of emissions changes the more constrained a county is by existing air quality standards. In equation (2.2), this equates to β_5-β_8 getting increasingly smaller as the regional design value approaches and passes the NAAQS threshold. These coefficients may even become negative, depending upon the magnitude of rebound in practice. We are particularly focused on comparing the *Near, Below* and *Near, Above* categories, since these counties are likely the most constrained.[7]

Beyond pollution impacts, we are ultimately interested in the effect of reductions in power plant emissions on human mortality, which drive the vast majority of co-benefit estimates in the regulatory impact assessments of emissions policies. Mirroring the analysis of how upwind emissions affect local $PM_{2.5}$ concentrations, we first assess how mortality varies with changes in local air quality driven by upwind emissions changes. However, rather than assess trends in month-to-month mortality, we focus instead on estimating the effect of the long-run change in local $PM_{2.5}$ from upwind sources on the long-run change in relevant mortality rates. We make this modification for two reasons. First, it is inherently difficult to estimate the causal effect of air quality on mortality rates due to numerous unobservable factors that also play a role in determining long-run health outcomes. Second, and more importantly, first-order mortality impacts from changes in air pollution exposure are generally the effect of years of exposure. Therefore, to assess the effect of the long-run change in air quality on the long-run change in mortality, we estimate

$$\left(\Delta\text{Mortality}\right)_c = \beta \left(\widehat{\Delta \text{PM}_{2.5}}\right)_c + \eta_{r(c)} + \varepsilon_c \tag{2.3}$$

where $\left(\Delta\text{Mortality}\right)_c$ and $\left(\widehat{\Delta \text{PM}_{2.5}}\right)_c$ are the changes in mortality and InMAP-estimated local $PM_{2.5}$ concentrations in county c, respectively, over the duration of each policy regime. For CAIR, we define this as the difference from 2009 to 2014, and for CSAPR/MATS from

[6] Or $PM_{2.5}$ precursors, which are the input to the InMAP model.

[7] The intuition is that the *Far, Below* counties have plenty of slack in their air quality constraint, and the *Far, Above* counties are almost certainly in non-attainment with NAAQS at baseline and presumably already working with EPA to reduce emissions, prior to any upwind shocks.

2015 to 2018.[8] In equation (2.3) we also include a NERC region fixed effect $\eta_{r(c)}$ to absorb level-based differences in mortality rates that are common across counties in the same NERC region and fixed over time.

We test our theory of regulatory rebound in the context of mortality by extending equation (2.3) similarly to the $PM_{2.5}$ case by allowing the effect of variation in upwind emissions to be heterogeneous based on the county's proximity to the $PM_{2.5}$ NAAQS in the baseline period before each policy is enacted. Specifically, we estimate

$$\begin{aligned}(\Delta\text{Mortality})_c = \\ &\beta_1 \left(\widehat{\Delta \text{PM}_{2.5}}\right)_c \times \mathbb{1}\{\text{Far, Below}\}_{Baseline} + \\ &\beta_2 \left(\widehat{\Delta \text{PM}_{2.5}}\right)_c \times \mathbb{1}\{\text{Near, Below}\}_{Baseline} + \\ &\beta_3 \left(\widehat{\Delta \text{PM}_{2.5}}\right)_c \times \mathbb{1}\{\text{Near, Above}\}_{Baseline} + \\ &\beta_4 \left(\widehat{\Delta \text{PM}_{2.5}}\right)_c \times \mathbb{1}\{\text{Far, Above}\}_{Baseline} + \\ &\eta_{r(c)} + u_c\end{aligned} \tag{2.4}$$

where the indicator variables for distance from the $PM_{2.5}$ NAAQS threshold are defined the same way as in equation (2.2). The outcome variable in the mortality regressions $(\Delta\text{Mortality})_c$ is calculated as the change in a relevant diagnosis-specific mortality rate per 100,000 residents in county c over the course of each policy regime.[9] Analogous to the PM case, evidence of regulatory rebound could be identified if β_1-β_4 are increasingly small, again with particular focus on comparing the two *Near* categories. This would be evidence that constrained localities *take back* some of the air quality (and ultimately, health-related) gains from policy-induced reductions in power plant emissions. Consistent with previous work to estimate the long-run mortality impacts of long-run changes in particulate matter exposure (e.g., Henneman, Mickley, and Zigler 2019), we estimate equation (2.4) by regressing the long difference in diagnosis-specific mortality rates on the long difference in InMAP-modeled $PM_{2.5}$ by OLS.[10]

[8] The CSAPR/MATS policy regime goes through 2019, however our mortality data from CDC currently end in 2018.

[9] Relevant diagnoses considered are those identified in previous work as being affected by poor air quality. See Section 2.5 for details.

[10] Related papers have instead estimated mortality impacts using count-based Poisson models, but we note that a Poisson model is not appropriate in our case given that the mortality changes we document are often negative. Also, there are very few zeroes in county-by-month mortality rates, even when limited to specific diagnosis categories.

2.5 Data Sources

2.5.1 Emissions and Plant Data

Our daily emissions data come from the US EPA's *Continuous Emissions Monitoring Systems* (CEMS).[11] For each day of operation for each electricity-generating unit (EGU), the CEMS data record heat input, total generation, CO_2 emissions, NO_x emissions, and SO_2 emissions. We merge these daily EGU emissions data with additional unit- and plant-specific features, including unit nameplate capacity, stack height, stack diameter (and area), geographic coordinates, NERC region, and dates of key events that may change units' emissions trajectories. These key events include retirements, fuel switching, and abatement-technology installations.[12]

2.5.2 EPA Monitor Readings, Design Values, and Distance Metrics

EPA maintains a national network of monitors that it uses to track local air quality and assess compliance with national air quality standards. These data are reported in EPA's Air Quality System and can be downloaded at the monitor-level for individual pollutants and temporal aggregations (*e.g.*, hourly, daily, annually, etc.) through EPA's data API. In this working paper, we focus on monitors with pollutant code 88101, which identify $PM_{2.5}$ monitors that meet the federal requirements to be used in NAAQS determinations, and we aggregate daily monitor readings to the monthly level by taking a simple average.

Individual $PM_{2.5}$ monitors are mapped to counties using their geographic coordinates, and we exclude any monitors outside of the continental U.S. Although most counties have just one $PM_{2.5}$ monitor that meets the standards for NAAQS calculations, some (often historically more-polluted) counties have multiple monitors. For instance, as of 2018, Cook County, Illinois and Los Angeles County had nine and ten monitors for $PM_{2.5}$, respectively, whereas San Francisco County has only one regulatory-grade monitor.

As discussed, determinations regarding local compliance with each pollutant-specific NAAQS are made on the basis of *design values* which are calculated using monitor readings from EPA's network. Design values are first computed at the monitor-level, and are based on pollutant- and duration-specific formulas which typically specify a data capture requirement and, upon regulatory approval, allow for exclusion of certain rare exceptional events that could affect concentration readings (e.g., wildfires, volcanic eruptions, etc.).[13] To calculate the design value for a given area, EPA selects the maximum of the valid monitor-level design values across all monitors within that area. Each year, EPA publishes monitor-level and

[11] We obtained this data through a partnership with EPA via its 2019 EmPOWER Air Data Challenge.

[12] While we currently input unmodified "raw" emissions into the InMAP model, we may in the future use these plant-level indicators to exclude EGUs whose emissions are likely not affected by regulations (*e.g.*, plants that we can identify as retiring due to trends in renewable energy or other reasons that are unrelated to the emissions policies we study).

[13] In our current analysis, in order to increase geographic coverage, we generally treat as valid the design values that are deemed "invalid" due to insufficient data capture, since they are published for public use annually and easily observable by local regulators. Versions of our analyses which exclude these invalid design values appear to reach the same general conclusions, but are less precise due to smaller sample size.

county-level design values, as well as design values for "previously designated nonattainment areas," which may be comprised of one or more counties. For example, the designated nonattainment area *Los Angeles-South Coast Air Basin, CA* includes Los Angeles, Orange, Riverside, and San Bernardino counties. Future iterations of this paper may instead consider geographic aggregations other than counties, which may correspond better to a local air district, such as a "commuting zone", as the spatial unit of analysis.

We currently define a county as "constrained" or not with respect to the NAAQS for annual $PM_{2.5}$, equal to 15.0 $\mu g/m^3$ through 2012 when it was reduced to 12.0 $\mu g/m^3$ in 2012 (with implementation in 2014). Specifically, we first calculate the average design value for a county in the baseline period of each policy (2007-2008 for the CAIR regime and 2013-2014 for the CSAPR/MATS regime) and calculate the distance between that value and the applicable NAAQS (15.0 and 12.0 for CAIR and CSAPR/MATS, respectively). We then assign each county to one of four groups based on this measure: $(-\infty, -1)$ (Far, Below), $[-1, 0)$ (Near, Below), $[0, 1]$ (Near, Above), and $(1, \infty)$ (Far, Above). In our estimation detailed below, we test for heterogeneity in regulatory rebound across these four mutually exclusive groups.

We construct our main estimation panel by matching county-level annual design values to $PM_{2.5}$ monitor readings at the specific site which determined the county-level design value in the given year (i.e., the monitor with the maximum DV in the county), which we call the "marginal" monitor site. This construction results in a panel in which each county only appears in the data once in each month of the sample. One alternative construction, which would include monitor readings from all monitors and assign distance based on each monitor's county-level design value, would over-represent counties which have multiple regulatory monitors. A third alternative would simply include all existing monitors and assign each monitor its own design value. However, that panel structure is less representative of reality on the ground, since individual monitor design values that are below the maximum in an area are not used for regulatory decisions. In any case, we run our empirical analysis using these two additional panel structures, but results are not meaningfully different from those using our preferred estimation panel.

2.5.3 Mortality Data

To assess long-run trends in mortality related to declining power plant emissions, we rely on Multiple Cause of Death (MCOD) data from the Centers for Disease Control and Prevention. This is a database of individual-level death records for the entire U.S. from 2005 to present. Each record contains demographic information and information about the event, including the primary cause of death identified by diagnostic codes contained in the International Classification of Diseases, Tenth Revision, Clinical Modification (ICD-10-CM). In addition to assessing trends in total mortality, we assess trends in several air quality-related conditions in order to be consistent with several epidemiological studies relied upon by EPA in estimating health co-benefits associated with proposed emissions policies. These conditions and associated ICD-10 code prefixes include the following: all cardiovascular diseases (I10.0-170.9), cardiopulmonary disease (I26-I28), ischemic heart diseases (I20-I25), lung and related cancers of the respiratory system (C33.0-C34.9), chronic lower respiratory diseases (J40.0-J47.0). We also examine trends in several other diseases not believed to be related

to air quality as a placebo check on our model. Mortality records are matched to local air quality at the county and month, which is the lowest level of spatio-temporal aggregation available in the confidential MCOD data.

2.6 Results and Discussion

In the results that follow, we report coefficient estimates from our estimating equations for our two policy regimes: CAIR and CSAPR/MATS. We begin by assessing how well $PM_{2.5}$ exposure predicted by InMAP correlates with observed local pollution levels. And then we test for regulatory rebound in both monitored $PM_{2.5}$ and mortality by allowing for heterogeneity in upwind (InMAP) $PM_{2.5}$ impacts by distance from the annual NAAQS threshold in the baseline period of each policy regime. For all regressions, we use a bandwidth of 1.0 $\mu g/m^3$ to assign the *Near, Above* and *Near, Below* categories. That is, counties with a design value in the baseline of the CAIR period (2007-2008) of between 14.0-15.0 $\mu g/m^3$ would be categorized as *Near, Below*, while counties with a design value of between 15.0-16.0 $\mu g/m^3$ are categorized as *Near, Above*. For the baseline of the CSAPR/MATS regime (2013-2014), we also use a bandwidth of 1.0 $\mu g/m^3$, but we note that distance metrics for that regime are determined with respect to the lower national threshold of 12.0 $\mu g/m^3$ that was enforced following the 2012 NAAQS update.

2.6.1 Local Air Pollution

Table 2.1 reports regression results from equations (2.1) and (2.2) for the CAIR regime. In column (1), we regress EPA $PM_{2.5}$ monitor readings (averaged monthly) on InMAP predictions for the grid cell in which each monitor is located, based upon all power plant emissions in the contiguous U.S. in the corresponding month. β in equation (2.1) gives the effect of *upwind*[14] emissions events on average monitor-level $PM_{2.5}$ concentrations at monitor i during month t in NERC region $r(i)$. If InMAP is accurately modeling pollution transport and $PM_{2.5}$ formation, we should expect to estimate a β parameter near one.

Column (1) of Table 2.1 suggests that the combination of aggregate EGU emissions and InMAP modeling slightly under-predicts (on average) the relationship between monitored $PM_{2.5}$ concentrations and precursor emissions originating at power plants. That is, β of 1.27 in equation (2.1) implies that every 1.0 $\mu g/m^3$ of modeled $PM_{2.5}$ is associated with 1.27 $\mu g/m^3$ of actual (monitored) $PM_{2.5}$ downwind. However, at the 5-percent level, although we reject that β is less than or equal to zero, we fail to reject that β differs from one (95% confidence interval of [0.87, 1.68]). One explanation for β being larger than one is that upwind $PM_{2.5}$ generation is correlated with local (near the monitor) $PM_{2.5}$-generating activity, thus creating some upward bias in the coefficient estimate.

The results of our preliminary empirical test for regulatory rebound under the CAIR regime are reported in column (2) of Table 2.1. Intuitively, we want to assess whether constrained localities respond to sustained reductions in upwind emissions by taking back the gap between local $PM_{2.5}$ concentrations and air-quality standards. In order to test this hypothesis, we allow the influence of upwind EGU emissions on local $PM_{2.5}$ concentrations to

[14]Upwind using InMAP's dispersion model.

differ by proximity to the NAAQS threshold for ambient $PM_{2.5}$ after the new emissions policy has been enacted. Since there are likely pre-existing regional differences in how power plant emissions affect downwind air quality – for example, due to natural differences in topology and/or average meteorology across regions – we control for heterogeneity in the pre-policy period (β_1 through β_4 in equation (2.2)). Our main focus is on the coefficients that are interacted with $CAIR$ (β_5 through β_8 in equation (2.2)), capturing heterogeneity during the CAIR policy regime (2009-2014). If regulatory rebound occurs, one would expect the estimated coefficients during the policy regime to be smaller the more constrained a county is, reflecting lower pass-through of local air quality benefits in areas that have relatively higher incentives to permit more economic activity following sustained drops in emissions from upwind sources.

The estimates in column (2) of Table 2.1 provide weak but suggestive evidence that regulatory rebound occurred in response to CAIR implementation. The first four estimates in column (2) suggest there exist some pre-policy differences in the impact of upwind emissions on local air quality, with the *Near* counties experiencing greater pass through of emissions than counties farther from the threshold in the pre-policy period. However, we cannot reject that β_1-β_4 are equal given that the 95% confidence intervals for the estimates are overlapping. Further, the coefficients on the *CAIR* interactions conform broadly with rebound, although they are imprecisely estimated. The coefficient on the *Near, Above* counties is relatively smaller than the coefficient on *Near, Below*, while the *Far, Above* coefficient is negative. This is what one would expect if areas with poorer ambient air quality at baseline face relatively larger incentives to permit more polluting activity when they experience a favorable external shock from upwind sources. However, although these estimates are potentially suggestive of rebound, they are relatively noisy. We note that there are many fewer localities classified as exceeding the annual standard in baseline period of CAIR (47 are classified as *Near, Above* and 25 as *Far, Above*, compared with 99 *Near, Below* and 667 *Far, Below*) which helps to explain the less precise coefficients on these CAIR-period interactions.[15] Overall, while these estimates do not provide irrefutable evidence of rebound, they also do not provide strong reason to reject it. Importantly, we note that while relying on InMAP to model downwind $PM_{2.5}$ formation offers a quick read on these potential relationships, its reduced complexity and static seasonal meteorology likely make it not the best suited model for this task. As we note, future iterations of this analysis will instead rely primarily on particle transport/dispersion predicted by the higher-complexity HYSPLIT model.

Table 2.2 reports empirical results for our second regime, CSAPR/MATS. The table structure is the same as in Table 2.1, but in this case we include interactions between the distance bins and the CAIR regime (2009-2014) in addition to the CSAPR/MATS regime (2015-2019), in order to allow distance-based heterogeneity to differ between the CAIR regime and the period prior to any policy implementation (2005-2008). We note that the distance metrics in this estimation sample are assigned to counties with respect to design values in the baseline period of the CSAPR/MATS regime (2013-2014) (rather than with respect to the CAIR baseline period of 2007-2008, as in Table 2.1).[16]

Similar to the CAIR analysis, InMAP under-predicts particulate pollution derived from

[15] We report results from a supplementary analysis that excludes California counties and collapses the *Near, Above* and *Far, Above* categories to a single *Above* category in Appendix Table B.3.

[16] A county's level of constrained-ness as of 2013-2014 seems to be the most relevant value to measure any

Table 2.1: The effect of upwind emissions on monitors' $PM_{2.5}$ levels under CAIR

Dependent variable: Monitor-average $PM_{2.5}$ (month of sample)

	(1) Pooled	**(2) Annual Standard**
$\widehat{PM_{2.5}}$	1.268***	
	(0.204)	
$(\widehat{PM_{2.5}})$ × Far, Below		1.093***
		(0.210)
$(\widehat{PM_{2.5}})$ × Near, Below		1.347***
		(0.209)
$(\widehat{PM_{2.5}})$ × Near, Above		1.442***
		(0.203)
$(\widehat{PM_{2.5}})$ × Far, Above		1.313***
		(0.269)
$(\widehat{PM_{2.5}})$ × Far, Below × CAIR		0.371
		(0.267)
$(\widehat{PM_{2.5}})$ × Near, Below × CAIR		0.476*
		(0.279)
$(\widehat{PM_{2.5}})$ × Near, Above × CAIR		0.287
		(0.262)
$(\widehat{PM_{2.5}})$ × Far, Above × CAIR		-0.491
		(0.350)
Fixed effects		
Monitor	Yes	Yes
Month by NERC region	Yes	Yes
Observations	62,451	62,451

*p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the monitor-month, limited to only one monitor per county (selected monitors represent the "marginal" monitor, defined as the monitor that determined the county's annual design value for $PM_{2.5}$ in the respective year). The dependent variable is the monthly average $PM_{2.5}$ concentration at the given monitor. $\widehat{PM_{2.5}}$ refers to InMAP-based estimates of the transport of $PM_{2.5}$ from upwind emissions at electricity-generating units to the monitor. Distance indicators denote the distance of the monitor's county from the Annual $PM_{2.5}$ NAAQS of 15.0 $\mu g/m^3$ in the base period of CAIR (2007-2008). Based on a bandwidth of 1.0 $\mu g/m^3$, there are 667 monitors categorized as *Far Below*, 99 as *Near Below*, 47 as *Near Above*, and 25 as *Far Above*. We two-way cluster errors at the same level as the specifications' fixed effects. The data sample for this set of regressions covers 2005-2015.

upwind emissions when we extend the panel to include the CSAPR/MATS years (2015-2019). At 95% confidence, we can reject a β of zero, although we fail to reject a value of one. To test for regulatory rebound during this second regime, we are focused primarily on the distance interactions with *CSAPR/MATS*, where evidence of rebound would be indicated again by smaller pass-through of upwind emissions in relatively more constrained counties. Of the four interactions of interest, the estimate for *Near, Above* is indeed relatively smaller than for *Near, Below* (and both smaller than *Far, Below*) but it is hard to interpret the increasingly large negative values. For example, the value for *Near, Above* counties of -1.4 implies that every 1.0 $\mu g/m^3$ decrease in $PM_{2.5}$ from upwind sources is associated with a 1.4 $\mu g/m^3$ *increase* in local $PM_{2.5}$. While this is theoretically consistent with rebound, the magnitude of the response seems extreme. Again, we note that InMAP is a highly simplified model of chemical transport, and the coefficients in column (1) of Tables 2.1 and 2.2 indicate that it does not accurately predict downwind $PM_{2.5}$, at least on average. We therefore emphasize that although weakly indicative of rebound, these InMAP-based estimates are preliminary. Future iterations of this paper will incorporate data-driven estimates of downwind $PM_{2.5}$ based instead on a state-of-the-art atmospheric dispersion model (HYSPLIT) which may generate different results.

2.6.2 Mortality

In this section, we report estimates of the effect of long-run changes in predicted $PM_{2.5}$ from upwind emissions on long-run changes in mortality rates for relevant conditions.[17] We note that mortality is a primary outcome of interest in our analysis given the large contribution of monetized mortality benefits to the projection of overall co-benefits under proposed emissions policies. We report estimates for five cardiac and respiratory health outcomes that are the focus of several studies that underlie EPA's mortality projections under recent policy proposals (e.g., Krewski et al. 2009 and Lepeule et al. 2012), in addition to all-cause mortality. These conditions include ischemic heart disease (IHD), all cardiovascular diseases, all cardiopulmonary diseases, lung cancer, and chronic obstructive pulmonary disorder (COPD).[18]

The estimates contained in Tables 2.3 and 2.4 are based on the regression presented in equation (2.4), and focus on long-run changes in predicted PM and mortality over the duration of the CAIR and CSAPR/MATS policy regimes, respectively. Each column represents a separate regression limited to mortality for the specified condition. Based on the granularity of the CDC mortality data, estimates are at the county-level, and represent the annual count of deaths per 100,000 county residents. To reduce noise in the data, we average over mortality rates (and InMAP-predicted $PM_{2.5}$) over two pre-policy years and two post-policy years for each regime. For CAIR, the long difference compares 2007-2008 values

changes against, since it determined the county's NAAQS compliance status just prior to the implementation of CSAPR and MATS. However, we recognize that 2013-2014 falls right at the end of the CAIR regime, and thus these CSAPR/MATS distance definitions may be confounded by any regulatory rebound that is occurring in response the earlier policy regime. In future work we will consider other measures of distance, e.g. the existence of offset permit markets.

[17] We report mortality estimates based on InMAP-modeled PM even though the results from the previous section indicated that InMAP may not be very well-suited for this task.

[18] Ischemic heart disease refers to heart problems caused by narrowed arteries around the heart. COPD includes chronic lower respiratory diseases, such as bronchitis, emphysema, and asthma.

Table 2.2: The effect of upwind emissions on monitors' $PM_{2.5}$ levels under CSAPR/MATS

Dependent variable: Monitor-average $PM_{2.5}$ (month of sample)		
	(1) Pooled	**(2) Annual Standard**
$\widehat{PM_{2.5}}$	1.278***	
	(0.176)	
$(\widehat{PM_{2.5}})$ × Far, Below		1.238***
		(0.192)
$(\widehat{PM_{2.5}})$ × Near, Below		1.189***
		(0.178)
$(\widehat{PM_{2.5}})$ × Near, Above		1.225***
		(0.237)
$(\widehat{PM_{2.5}})$ × Far, Above		1.591***
		(0.393)
$(\widehat{PM_{2.5}})$ × Far, Below × CAIR		0.239
		(0.265)
$(\widehat{PM_{2.5}})$ × Near, Below × CAIR		0.216
		(0.272)
$(\widehat{PM_{2.5}})$ × Near, Above × CAIR		0.135
		(0.289)
$(\widehat{PM_{2.5}})$ × Far, Above × CAIR		0.097
		(0.578)
$(\widehat{PM_{2.5}})$ × Far, Below × CSAPR/MATS		-0.203
		(0.573)
$(\widehat{PM_{2.5}})$ × Near, Below × CSAPR/MATS		-0.529
		(0.544)
$(\widehat{PM_{2.5}})$ × Near, Above × CSAPR/MATS		-1.418**
		(0.716)
$(\widehat{PM_{2.5}})$ × Far, Above × CSAPR/MATS		1.066
		(2.118)
Fixed effects		
Monitor	Yes	Yes
Month by NERC region	Yes	Yes
Observations	86,267	86,267

*p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the monitor-month, limited to only one monitor per county (selected monitors represent the "marginal" monitor, defined as the monitor that determined the county's annual design value for $PM_{2.5}$ in the respective year). The outcome variable is the monthly average $PM_{2.5}$ concentration at the given monitor. $\widehat{PM_{2.5}}$ refers to InMAP-based estimates of the transport of $PM_{2.5}$ from upwind emissions at electricity-generating units to the monitor. Distance indicators denote the distance of the monitor's county from the Annual $PM_{2.5}$ NAAQS of 12.0 $\mu g/m^3$ in the base period of CSAPR/MATS (2013-2014). Based on a bandwidth of 1.0 $\mu g/m^3$, there are 807 monitors categorized as *Far Below*, 72 as *Near Below*, 30 as *Near Above*, and 22 as *Far Above*. We two-way cluster errors at the same level as the specifications' fixed effects. The data sample for this set of regressions covers 2005-2019.

to 2013-2014 values, and for CSAPR/MATS the comparison is 2013-2014 to 2017-2018. We note that the results presented currently are essentially "reduced form" estimates, based on evidence from the air pollution regressions that InMAP $PM_{2.5}$ is not perfectly correlated with actual (monitored) $PM_{2.5}$. In future work, we hope to move beyond the reduced form to generate estimates of mortality impacts that can be compared directly to projections contained in regulatory impact assessments of emissions policies. Positive estimates for each of the mortality regressions would indicate that $PM_{2.5}$ and relevant mortality move in the same direction, which one would expect if pollution is bad for health. Analogous to the previous section, evidence of regulatory rebound would be indicated if the coefficients on the distance-based interactions become relatively smaller the farther a county moves out of attainment with the prevailing NAAQS. This would suggest lower pass-through of expected mortality benefits from CAIR or CSAPR/MATS in counties that were constrained *ex ante* by existing air quality standards at the time that each emissions policy was enacted.

The coefficient estimates in Table 2.3 do not provide strong evidence of rebound in mortality under the CAIR period, which is not totally surprising given the imprecision of the first-stage estimates in Table 2.1. The presence of several negative coefficients – indicating higher mortality as pollution levels fall – is unexpected, although the estimates are generally imprecise and very small in magnitude (especially when interpreted as an annual rate per 100,000 county residents). The results in Table 2.4 for the CSAPR/MATS regime also do not provide strong evidence of mortality-related regulatory rebound. However, we again note that these estimates are preliminary, and are based on InMAP-modeled $PM_{2.5}$ which we show is not a perfect predictor of actual $PM_{2.5}$. We also note that there are far fewer counties that are assigned to the *Above* categories than the *Below* categories, which helps to explain the relatively noisier coefficients on these interactions. In future work, we may incorporate other measures of a county's level of constrained-ness with respect to NAAQS, including possibly the existence of local markets for emissions offset permits.

2.6.3 Robustness

Based on our primary estimates, we have identified weak but suggestive evidence of regulatory rebound in local pollution levels (although not in long-run mortality). However, since a large portion of counties assigned to the *Far, Above* category are located in California – a state which was only marginally impacted by the CAIR and CSAPR/MATS policy changes – we replicate our analyses with all California counties excluded from the sample. These supplementary results are reported in Appendix Section B.1. When excluding California counties, the CAIR regime estimates reported in Appendix Table B.1 continue to indicate weakly suggestive evidence of rebound based on our proposed empirical test. Similarly, estimates for the CSAPR/MATS regime, in Appendix Table B.2, offer limited evidence of rebound. We also note that once excluding California counties from our analysis under the CSAPR/MATS regime, we are left with only two non-California counties in the *Far, Above* category. Therefore we conduct a final set of CSAPR/MATS regressions in Appendix Table B.3 in which we combine *Near, Above* and *Far, Above* to a single *Above* category. In this case, we again find some evidence of rebound, although the majority of estimates are not statistically significant. Overall, we are encouraged by the limited evidence we have obtained of rebound occurring in practice, but continue to interpret these InMAP-based estimates with

Table 2.3: The effect of changes in upwind emissions on changes in county mortality under CAIR

	(1) All	(2) IHD	(3) Cardiovasc.	(4) Cardiopulm.	(5) Lung	(6) COPD
$(\widehat{\Delta \mathrm{PM}_{2.5}})$ × Far, Below	1.390	2.667	−3.448	−1.025	−1.899	−1.210
	(29.161)	(6.205)	(10.233)	(0.634)	(2.615)	(2.877)
$(\widehat{\Delta \mathrm{PM}_{2.5}})$ × Near, Below	−10.467	1.733	−5.301	−1.489**	−3.461	−2.345
	(27.330)	(5.815)	(9.590)	(0.594)	(2.451)	(2.696)
$(\widehat{\Delta \mathrm{PM}_{2.5}})$ × Near, Above	13.694	7.507	7.614	−0.660	−1.195	2.100
	(29.027)	(6.177)	(10.186)	(0.631)	(2.603)	(2.864)
$(\widehat{\Delta \mathrm{PM}_{2.5}})$ × Far, Above	4.782	5.681	0.383	−1.335	−0.091	1.953
	(48.544)	(10.330)	(17.035)	(1.056)	(4.353)	(4.789)
Fixed effects						
NERC Region	Yes	Yes	Yes	Yes	Yes	Yes
Observations	438	438	438	438	438	438

*p<0.1; **p<0.05; ***p<0.01. The outcome variables is the change in the county-level mortality rate per 100,000 residents for the stated health condition from the pre-period to the post-period of the policy. For CAIR, this change is calculated as the difference between the 2013-2014 average rate and the 2007-2008 average rate. $\widehat{\Delta PM_{2.5}}$ refers to the change in average InMAP-based exposure from 2007-2008 to 2013-2014. Distance indicators denote the distance of the county from the Annual $PM_{2.5}$ NAAQS of 15.0 $\mu g/m^3$ in the base period of CAIR (2007-2008). We cluster errors by NERC region.

Table 2.4: The effect of changes in upwind emissions on changes in county mortality under CSAPR/MATS

	(1) All	(2) IHD	(3) Cardiovasc.	(4) Cardiopulm.	(5) Lung	(6) COPD
$(\widehat{\Delta \text{PM}_{2.5}})$ × Far, Below	4.371	4.412	22.702	0.229	8.552	−2.114
	(51.588)	(12.371)	(17.032)	(1.358)	(6.134)	(6.307)
$(\widehat{\Delta \text{PM}_{2.5}})$ × Near, Below	30.284	4.818	25.392	−2.084	10.678	1.612
	(59.278)	(14.215)	(19.570)	(1.560)	(7.049)	(7.247)
$(\widehat{\Delta \text{PM}_{2.5}})$ × Near, Above	−2.787	10.729	7.210	−0.577	15.703	2.038
	(90.715)	(21.754)	(29.949)	(2.388)	(10.787)	(11.090)
$(\widehat{\Delta \text{PM}_{2.5}})$ × Far, Above	3.733	13.569	23.931	−0.041	7.279	0.175
	(168.430)	(40.390)	(55.606)	(4.434)	(20.028)	(20.591)
Fixed effects						
NERC Region	Yes	Yes	Yes	Yes	Yes	Yes
Observations	437	437	437	437	437	437

*p<0.1; **p<0.05; ***p<0.01. The dependent variable is the change in the annual county-level mortality rate per 100,000 residents for the stated health condition from the pre-period to the post-period of the policy. For CSAPR/MATS, this change is calculated as the difference between the 2017-2018 average rate and the 2013-2014 average rate per 100,000 residents. $\widehat{\Delta PM_{2.5}}$ refers to the change in average InMAP-predicted exposure from 2013-2014 to 2017-2018. Distance indicators denote the distance of the county from the Annual $\text{PM}_{2.5}$ NAAQS of 12.0 $\mu g/m^3$ in the base period of CSAPR/MATS (2013-2014). We cluster errors by NERC region.

some caution.

We also conduct several placebo tests related to the mortality analysis, by estimating the impact of changes in predicted $PM_{2.5}$ on changes in long-run mortality for several conditions that are not known to be associated with air quality. In general, we find no effects, as expected. These estimates are reported in Appendix Tables B.4 and B.5.

2.7 Conclusion

This paper begins to investigate co-benefits accounting in theory and practice. We focus on two distinct policy regimes, the CAIR regime (2009-2014) and the CSAPR/MATS regime (2015-2019), during which the U.S. saw unprecedented reductions in the emissions of harmful criteria pollutants associated with power generation. We highlight the potential importance of policy interactions in the determination of health co-benefits associated with reductions in these pollutants. We introduce an empirical strategy for evaluating the indirect impacts of power plant emissions on air quality in down-wind locations that accounts for this "regulatory rebound", which we theorize may manifest in ex ante constrained localities. To approximate the transport of precursor pollution from power plants and formation of $PM_{2.5}$, we rely on the InMAP model, a reduced-form chemical transport model that allows us to get a preliminary read on how emissions changes affect down-wind pollution levels, and ultimately, long-run mortality.

Based on our current empirical strategy and the projections of downwind $PM_{2.5}$ delivered by the InMAP model, we currently find weakly suggestive evidence of regulatory rebound in practice. While our proposed empirical test does not deliver irrefutable evidence of rebound in local pollution levels or long-run health impacts in constrained areas, we also cannot rule it out. We note that there are several important limitations to our current research design, which we aim to improve upon in future iterations of this work. First, while useful for obtaining preliminary suggestive estimates, the InMAP model may not be the best-suited tool for this task. The effects we hope to detect require accurate, precise estimates of how emissions particles travel through space and time, and InMAP – which abstracts away from year-to-year variation in meteorology – may misattribute emissions to down-wind locations. Indeed, our efforts to validate the InMAP projections in our baseline model (column (1) in Tables 2.1 and 2.2) indicate that InMAP does not perfectly predict local pollution. As we mention, future work will incorporate more precise particle trajectory modeling based instead upon the higher-resolution and more computationally-intensive HYSPLIT model. Estimates of local pollution levels based on emissions particle trajectories projected by HYSPLIT will support a more rigorous test of our regulatory rebound hypothesis, and allow us to directly compare our estimates of benefits to their projected levels reported in EPA's prospective regulatory analyses. In addition to incorporating more complex pollution transport modeling, we will also continue to refine our empirical approach. We may consider other relevant spatial units beyond individual counties, as well as broaden our definition of 'constrained' to include qualitative indicators, such as the existence of offset permit markets (as explored by Shapiro and Walker 2020). Overall, we aim to inform the current debate on how co-benefits should be properly accounted for, and provide rigorous empirical evidence on how reductions in power plant pollution impact local air quality and associated health outcomes.

Chapter 3

Income-Related Heterogeneity in Extreme Temperature and Health

3.1 Introduction

Whether and to what degree humans are harmed by a changing climate is currently an active area of research in empirical environmental economics. There is evidence that extreme temperatures and the unpredictable weather events made more likely by climate change (such as tropical cyclones) affect a wide range of economic outcomes, summarized in detail by Carleton and Hsiang 2016. Examples of potential negative impacts of climate change include reductions in agricultural yields (Aufhammer, Ramanathan, and Vincent 2012 and Schlenker and Roberts 2009), deceleration of macroeconomic growth (Burke, Hsiang, and Miguel 2015), strained energy systems and infrastructure (e.g., Jaglom et al. 2014), increased civil conflict and interpersonal violence (Burke et al. 2009 and Ranson 2014), reduced labor supply and productivity (Graff-Zivin, Hsiang, and Neidell 2018), increased mortality (Deschenes and Greenstone 2011), and greater disease prevalence (e.g., Burke, Hsiang, and Miguel 2015).

Of particular importance when projecting climate damages into the future is the range of possible actions that individuals might take to moderate harmful impacts, i.e. human adaptation to climate change. Much attention has been focused in recent years on envrionmental policy measures that might mitigate the accumulation of harmful gases in the atmosphere, such as taxation of carbon emissions and increasingly stringent regulations for other pollutants. More recently, researchers have focused on the adaptive measures that individuals and other agents can take to adapt to actual or expected climate. Example adaptive behaviors include investments in residential cooling systems, strategic urban planning, or improved public health services, such as early warning systems. Also important to the estimation of climate change-related damages is understanding whether and how populations with different underlying characteristics are differentially affected by the same climate conditions, since this heterogeneity could have important implications for the distribution of climate damages (or benefits) and thus the design of optimal climate policy.

In the United States, public assistance provided by the government to low-income populations, administered through in-kind benefit programs like Medicaid or through direct cash welfare via programs like the Supplemenal Nutrition Assistance Program, is aimed broadly

at redistributing resources from the rich to the poor. However, one unintended benefit of such welfare programs could also be improved protection from climate-related shocks, either through directly increasing the baseline health stock of program recipients (via improved nutrition or investments in preventive health care) or indirectly through loosening the budget constraint to allow investment in other defensive measures or in the general stock of health.

In this chapter, I return to the potential for public health insurance (e.g., Medicaid) to be used as a tool to mitigate the health impacts of higher temperatures, rather than air pollution exposure as studied at length in Chapter 1. I begin by replicating previous work on the health impacts of extreme temperature using Texas as a case study. I confirm the "U"-shaped relationship documented by previous researchers, where both extreme cold and extreme hot temperatures are associated with higher morbidity, relative to a mild day. Furthermore, although as of 2020 Texas has not expanded its Medicaid program under the ACA, I provide descriptive (non-causal) evidence that adult recipients of Medicaid may be relatively less vulnerable to extreme climate shocks. I supplement this with several additional analyses in which insurance status is replaced by regional income measures, in an attempt to disentangle the "Medicaid effect" from the "income effect". Whether public health insurance can be used to counteract the effects of extreme weather is a potentially policy-relevant question, particularly as climate change is making extreme temperatures and related climate events more common over time. If other studies are identifying problems that can be (or already are) solved by public assistance programs like Medicaid, then the policy recommendations would be clear.

Using administrative patient-level data on hospital use and insurance status, I estimate that health care utilization for temperature-related illnesses is relatively lower among the population of patients insured by Medicaid, but only when limited to the subset of patients 18 years and older. For children, I find the opposite result: Medicaid recipiency is associated with higher morbidity on extreme temperature days. Without more detailed information on patient-level characteristics, it is difficult to ascertain the reasons for this difference, but it might suggest that the positive health benefits of Medicaid recipiency strengthen over time, and that the protective Medicaid effect should theoretically be lower for children simply due to their young age (and fewer number of potential years enrolled in the program). Alternatively, the differences I estimate may be the result of unobservable selection bias which I cannot control for with the current research design. Indeed, I emphasize that, unlike the air pollution and Medicaid analysis summarized in Chapter 1, in this chapter I am not leveraging quasi-random variation in insurance coverage and so these estimates should not be interpreted as identifying a causal effect of that coverage. However, my data sample includes only patients that were ultimately admitted to the hospital following an initial ER visit, so it seems reasonable to assume that these patients are indeed acutely ill, which may alleviate some concerns about insurance-related moral hazard that has been documented by previous researchers. As a further check on this issue, in some specifications I focus on a set of more severe temperature-related diagnoses for which the demand for health care is likely to be (at least almost) universal, regardless of insurance status.

This study fits loosely into several existing literatures in economics and related fields. First, it contributes to the existing literature on the health impacts of climate, providing some of the first evidence (to my knowledge) that extreme temperatures lead to increased

morbidity among patients with any form of insurance.[1] Second, based on the potentially strong assumption that Medicaid recipiency is a reasonable proxy for an individual having relatively lower income, this study adds to a well-developed literature documenting how income is related the utilization of and demand for health care services (e.g., one influential example is Hall and Jones 2007, who find large income elasticities in the demand for health and health care). Third, this study is related to an emerging literature that uses plausibly exogenous variation in local climate to estimate health impacts among individuals receiving public assistance of various forms, or to quantify the added value of such programs in the wake of extreme climate events (e.g., existing work in this area includes Heutel, Miller, and Molitor 2020, Bishop, Ketcham, and Kuminoff 2018, and Deryugina 2017).

The chapter proceeds as follows. Section 3.2 provides an overview on the state of the literature on the relationship between temperature and health, presents institutional details about the Medicaid system that are relevant to this analysis, describes a conceptual framework of how Medicaid recipiency and income may modify the temperature-health relationship, and discusses current knowledge on income-related hetereogeneity in the demand for health care services. Section 3.3 provides information on the data sources used for estimation of morbidity impacts, including the Texas hospital data, and Section 3.4 reports relevant summary statistics. Sections 3.5 and 3.6 detail my econometric strategy and empirical results. Section 3.7 discusses the implications of these results, and concludes.

3.2 Background

3.2.1 Extreme Temperature and Human Health

The human body's thermo-reguatory system allows for continued functioning during ambient exposure to extreme hot or cold temperatures. However, despite the body's natural self-regulating behaviors, there are a variety of mechanisms through which extreme temperatures can lead to compromised health. Low temperatures cause the body's blood vessels and arteries to narrow, restricting blood flow and reducing oxygen to the heart. The heart must then pump harder to circulate blood through the constricted blood vessels, which increases blood pressure and heart rate. These physiological responses, particularly when coupled with physical exertion, can lead to serious medical conditions such as heart attack and stroke. Conversely, blood vessels will dilate in extreme heat, allowing more blood to flow to the surface of the skin, where heat can be lost to the air. With more blood near the surface of the skin, however, less is available to the heart, and heart rate will increase to maintain cardiac output. During this process, chemicals are released that make blood stickier and more likely to clot, and excess sweating triggers the loss of electrolytes that are critical to the functioning of both the heart and nervous system. These responses can lead to conditions such as heat exhaustion and heat stroke, which are sometimes fatal (and are particularly dangeous for the young and the elderly, since their thermo-regulatory systems are relatively less advanced).

[1] While there exist numerous epidemiological studies on temperature-related morbidity, previous economics studies, such as Deschenes and Greenstone 2011, have focused primarily on *mortality* related to extreme weather and/or temperature.

In one of the early economic studies on the relationship between temperature and health, Deschenes and Greenstone 2011 estimate a statistically significant association between daily mean temperatures and the annual mortality rate in the United States, with extreme temperatures at both ends of the spectrum associated with elevated risk of death. Specifically, using inter-annual variation in daily temperatures and aggregate state-by-year mortality data, the authors find that an additional day with a mean temperature above 90°F, relative to a day in the 50°-60°F range, leads to a 0.11 percent increase in the annual mortality rate. Similarly, a day below 20°F is associated with an increased mortality rate of about 0.08 percent. Notably, the authors uncover significant heterogeneity in the responses to extreme temperatures across age groups, with the elderly and the young being relatively more harmed.

While mortality is an important indicator of population health as a whole, as well as the largest cost when monetized using values for statistical life, focusing on this measure alone is likely to understate health damages associated with severe weather. In particular, as a result of exposure to extreme temperatures, individuals may suffer damages to health that don't immediately lead to loss of life, but should arguably be captured in a comprehensive analysis of extreme temperatures and health. As summarized in Chapter 1, there has been recent work exploring the morbidity impacts of air pollution exposure in the U.S. (e.g., Deryugina et al. 2019 and Schlenker and Walker 2016), but morbidity has been relatively less explored by economists working on issues related to climate. Instead, the latest work on climate and health has sought to investigate the adaptive behaviors and investments that individuals and societies can make in an effort to mitigate the harmful effects of exposure to extreme climate conditions. As discussed, these adaptative measures should be accounted for when estimating future damages of climate change since they will temper the associated health-related reductions in overall welfare. Motivated by the striking decline in U.S. mortality over the twentieth century, Barreca et al. 2016 provides some of the clearest evidence that adaptation is worth measuring. The authors assemble a comprehensive database of historical U.S. death records and daily temperatures, as well as measures of adaptation that might modify the temperature-mortality relationship. They find that temperature-related mortality has declined by 75% since 1960, a trend which they link primarily to the diffusion of residential air conditioning.

Another recent study on heterogeneity in the temperature-mortality relationship investigates the role of historical adaptation to climate. Heutel, Miller, and Molitor 2020 bring together data on the universe of Medicare beneficiaries and daily weather readings and find that the mortality effects of extreme temperatures vary widely depending on average local climate. The authors assign U.S. ZIP codes to one of three distinct groups based on historical average temperature (e.g., cold, moderate, or warm ZIPs) and find that the marginal effect of an extremely hot or extremely cold day varies dramatically across these categories. For example, they find that the mortality effects of hot days are much larger in cool regions (increased mortality of 3.3 deaths per 100,000 residents for a day hotter than 95°F, relative to a moderate 65°-70°F day) versus slightly negative in the hottest third of ZIP codes (a reduction in mortality of 0.1 deaths per 100,000 people). Similarly, the authors find that extremely cold days are much less harmful in cool regions than in the warmest regions: an increase in 0.2 deaths per 100,000 versus 2.4, respectively. These findings have important implications for projecting climate-related mortality into the future, if climate adaptation is indeed an important modifier of the temperature-mortality relationship. The authors sug-

gest that incorporating this regional heterogeneity and adaptation to local climate over time leads to predictions that atmospheric warming will actually *decrease* mortality in the coldest third of the U.S. while increasing mortality in the warmest third.

One dimension of heterogeneity and adaptation that is of increasing interest in the study of temperature and mortality/health is income. In a recent study, Carleton et al. 2020 estimate age-specific causal relationships between climate and mortality based on data that covers 40% of the global population, treating income as well as average climate (as in Heutel, Miller, and Molitor 2020) as the primary determinants of adaptation. They find that relative to a day at 20°C, a day at 35°C increases the global mortality rate by 2.94 deaths per 100,000. Their analysis suggests that by year 2100, this rate will be reduced to only 0.89 deaths per 100,000, with income accounting for 88% of the reduction. The finding that income is adaptive at the hot end of the temperature distribution provides the primary motivation for examining the influence of income at a more granular level in the present study. While the wide temporal and spatial scope of the causal estimates presented in the Carleton et al. 2020 study are undoubtedly useful for generating more precise estimates of the social cost of carbon that account for the costs and benefits of adaptation, the authors rely on highly aggregated income measures to arrive at their headline results. In this chapter, I exploit information on individual-level income (which I infer via Medicaid enrollment status) to understand how income might influence temperature-health impacts on an individual-level. Again, however, I emphasize that the Medicaid and income-related estimates in this chapter are in general not causal, and may be biased by unobservables correlated with both Medicaid status and propensity to seek emergency care in response to extreme temperatures. Still, to my knowledge, this study is a first step in examining income-related heterogeneity in climate impacts for individuals.

3.2.2 Medicaid Institutional Details

This section expands upon some of the basic background information about the Medicaid program that was first introduced in Chapter 1. I also provide additional details about the program in Texas, a state which unlike California did not adopt an expansion under the Affordable Care Act. Medicaid is the third largest U.S. entitlement program after Social Security and Medicare, with approximately $500 billion in annual spending in recent fiscal years. The program was established alongside Medicare in 1965 through amendments to the Social Security Act, and underwent significant expansions in the 1980s and 1990s, as well as more recently in 2012 with the ACA. Medicaid provides counter-cyclical social insurance, in that its enrollment expands to meet rising needs during an economic downturn, when employment rates fall and people lose employer-based health coverage. During the recession of 2008 and its aftermath, Medicaid enrollment increased by approximately 10 million people, more than half of which were children. As of 2015, the program served nearly 100 million low-income Americans, with children under age 20 accounting for over 40% of enrollees during the year. Importantly, Medicaid is an uncapped entitlement program; the government is not permitted to limit the number of eligible people who can enroll, and it must pay for any covered services for enrollees.

Eligibility for the program is based generally on income and family size, and can also be granted automatically by eligibility for certain cash welfare programs. The fiscal burden

of Medicaid is jointly shared by U.S. states and the federal government, with the federal payment share inversely proportional to the state's per capita income (though the federal share is always greater than or equal to 50% of costs). Federal guidelines for the program are quite broad and the states have historically had significant discretion in designing and administering their programs–states independently determine the amount, duration, and scope of the services they provide under Medicaid (though the services must meet a minimum standard). For example, states must cover hospital and physician services, but they can limit the number of hospital days or physician visits they pay for. As a result of this flexibility, Medicaid benefits packages can and often do vary widely from state to state.

As of 2015, approximately 16% low income Texans were insured by Medicaid, and Medicaid payments to providers represented over 25% of the state's budget in FY 2013. Yet, according to healthinsurance.org, Texas is tied with Alabama for having the most stringent eligibility criteria in the nation. In order to qualify for Medicaid coverage in Texas (if not already granted automatic coverage through another federal assistance program) , a family of four, for example, must not have annual income before taxes exceeding $48,708. Further, the state government of Texas has not accepted Medicaid expansion under the ACA, which if it were accepted would cover an additional 2.5 million Texans.[2] Regardless of whether Texas adopts a Medicaid expansion in the future, it remains the case that it is a program whose enrollment and costs are not capped. If the poor are relatively more damaged by a changing climate and are less able to adapt than those with higher incomes, federal and state expenditures on Medicaid could rise exponentially to meet increased demand. Alternatively, if it can be shown that public assistance programs like Medicaid provide (unanticipated) supplemental health benefits to vulnerable populations that suffer increasingly extreme temperatures, then it might suggest that the expansion of these programs could serve as another line of defense in the broader fight against climate change.

There has been research demonstrating that Medicaid recipients may use health care differently than individuals with other types of insurance (see, for example, Taubman et al. 2014). It is also the case that not all providers must accept Medicaid insurance, which could result in unobservable supply-side influences on observable health care utilization that differ between Medicaid and non-Medicaid patients. However, following the federal Emergency Medical Treatment & Labor Act (EMTALA) of 1986, all acute care facilities that have emergency rooms and treat Medicare patients (which constitutes essentially all public and private hospitals in the U.S.) cannot turn away patients regardless of ability to pay. Hence, for "true" emergencies, Medicaid patients should not be seen at emergency rooms at significantly lower rates than patients with other forms of insurance. This suggests that comparing temperature-related health outcomes across Medicaid/non-Medicaid groups – at least for relevant conditions with sufficient severity – might still be valuable even in the absence of exogenous variation in insurance coverage. I return to this issue in the empirical analysis that follows.

[2]Statistics reported at https://www.healthinsurance.org/texas-medicaid/.

3.2.3 Health Insurance-Related Mitigation of Climate Damages

The potential reasons that Medicaid might help individuals adapt to extreme climate events are similar to those discussed with respect to poor air quality in Chapter 1. I therefore just briefly summarize the intuition here. Enrollment in the Medicaid insurance program might lead to improved health outcomes (both in general and as related to extreme climate events) through two main channels. First, a substitution effect leads to increased utilization of health care services since they are now less costly. In fact, upon enrollment in Medicaid, many health care services (including preventive and emergency visits) are provided completely free of charge. As with air quality-related illness, this could result in recipients seeking more preventive health care in earlier periods that could improve baseline health and possibly increase resiliency to future health shocks. An income effect, arising from a loosening of the budget constraint since health care now has an effective price of zero, results in more disposable income available for the consumption of other non-health care goods. This additional income might be shifted toward food consumption (also potentially increasing individuals' baseline stock of health) or could be used to invest in durable goods or technologies that allow individuals to better insulate themselves from extreme weather conditions. One potential investment that could be undertaken that has been shown to be particularly valuable in reducing temperature-related mortality is residential air conditioning, as documented by Barreca et al. 2016. Another potential investment is personal transportation. Low-income households are generally more frequent consumers of public transportation services, which implies potentially more time spent outside exposed to ambient temperatures. Extra household resources might be spent toward the purchase of an automobile, reducing the amount of time that household members spend outside. Such large durables investments are often quite costly (for example, a residential A/C unit costs on average $5,327 in the U.S. in 2018)[3] but could be made more likely by an effective redirection of income away from health spending upon enrollment in public insurance.

3.2.4 Income Effects vs. "Medicaid Effects"

As discussed, one potential issue with using hospital admissions data as a proxy for individuals' true health status is that there are likely important differences in the demand for health care across the income distribution, both in general as well as during periods of extreme temperature. This could result in non-random selection into the Texas hospital visit sample. Indeed, Hall and Jones 2007 find an income elasticity for health care spending above one, suggesting that people spend an increasing fraction of income on health care as income increases. This behavior is explained by the authors as the fact that the marginal utility of an additional year of health does not decline in the same way that the marginal utility for normal consumption goods does. On the other hand, among the poor, enrollment into Medicaid might change recipients' demand for health care services as compared to the non-Medicaid poor or patients with other forms of insurance, again resulting in a non-randomly selected sample of patients that present to the ER for any medical issue. This has been shown to be true empirically. Results from an experiment on Oregon's Medicaid program that randomly granted Medicaid status to a sample of the eligible population revealed that

[3]Estimate obtained at https://www.homeadvisor.com/cost/heating-and-cooling/install-an-ac-unit/.

Medicaid patients utilize emergency departments for all forms of care at higher rates than non-Medicaid patients, including for conditions that may be better treated in a primary care setting (Taubman et al. 2014). One hypothesized reason for this result is that visits are essentially free under Medicaid, whereas ER co-pays for patients with other forms of insurance are often fairly expensive, which encourages relatively higher utilization of primary care rather than emergency care among non-Medicaid populations.

For these reasons, we should expect health care utilization to differ across both the income distribution as well as between Medicaid/non-Medicaid enrollees conditional on being low income, even in the absence of any exogenous health shocks. What is less clear (and unfortunately not empircally testable with the current research design) is whether the health care demand response of each of these groups is heterogeneous with respect to the short-run temperature variation that I rely on for identification in this study. As a hypothetical example, it could be the case empirically that adults with higher incomes are relatively more likely to seek immediate medical attention when they fall ill during times of extreme heat or cold, which could lead to a higher observed hospital visit rate, regardless of the true relative physiological health responses of the different groups. For children, non-random selection could also be a concern, depending on the relative incentives and constraints that parents/guardians face when deciding whether to seek medical care for a sick child. Previous work has shown that the poor obtain a relatively higher marignal return of investing in children's health, which could lead to higher observed rates of health care utilization among poor children in response to exogenous health shocks. However, the income effect identified by Hall and Jones 2007 would work in the opposite direction: one should expect to see higher visit rates among *higher income* children in response to the same shocks, simply because higher income households spend relatively more on health care services.

The purpose of this discussion is to highlight that while a descriptive analysis of differences in temperature-related health care utilization across Medicaid/non-Medicaid individuals might be valuable (especially if limited to sufficiently severe conditions for which health care demand should be nearly universal), there may still be significant selection bias operating in the background. I emphasize that the results in this chapter are meant to be interpreted as suggestive rather than causal, unlike the Medicaid-related estimates in Chapter 1.[4] Still, I conduct several different versions of my empirical analysis, testing models with different proxies for income. First, in some specifications I replace the indicator for an individual's Medicaid status with a measure of per capita annual income in the individual's home county. I additionally test specifications where the individual's Medicaid status is instrumented with the measure of county-level per capita income. It is possible that county-level income is uncorrelated with any unobservables that influence likelihood of inpatient admission, except through its effect on Medicaid status. I also run these same specifications but instead replace the per capita income variable with a measure of county-level poverty (although since income and povety are closely related, these models generate very similar results, hence I only report results from the models based on income). As a final check on this issue, I also report results when the sample is limited to visits for a set of relatively

[4]Unfortunately, given that Texas has not expanded the Medicaid program under the ACA, it is not possible to conduct a similar analysis leveraging spatial differences in potential policy exposure, as in Chapter 1.

severe diagnoses for which the demand for healthcare should be nearly universal, e.g. cardiac arrest and congestive heart failure.

3.3 Data Sources

3.3.1 Health Data

The hospital utilization data is collected by the Texas Department of State Health Services (T-DSHS), and represents a complete universe of all hospital visits in the state of Texas over the study period (2001-2010). The dataset includes basic patient information (month of visit, age group, sex, race, ethnicity, and home ZIP code) as well as medical information, including the principal cause of admission and up to 20 additional diagnosis codes, whether the admission was planned or unplanned, the route into the hospital (emergency versus scheduled), the patientâĂŹs primary health insurance provider (private insurance, Medicare, Medicaid, etc.), the patient's length of stay in the hospital, and a list of all procedures performed. There are also coded indicators for the patient's risk status and severity of conditions during their hospital stay. To protect patient confidentiality, these records are aggregated to the calendar quarter (3-month) level. To my knowledge, data like these have not been previously exploited by economists to explore the health impacts resulting from extreme climate conditions. Patient data is linked to the weather data (described below) using the patient's home ZIP code. For computational reasons, these data are then aggregated to the county-level, by taking a population-weighted average of weather variables across ZIP codes in each county. Note that this processing implicitly assumes that patients are exposed only to the weather conditions in their home county.

Following the epidemiology literature and previous studies within the economics field on the temperature-health relationship, I focus on admission rates for respiratory and cardiovascular diagnoses that might be induced or exacerbated by extreme weather conditions.[5] Using the ICD diagnostic classification system described in Chapter 1, I limit to respiratory diagnosis codes with the following prefices: 491 (bronchitis), 492 (emphysema), 493 (asthma), 496 (chronic obstructive pulmonary disease) for adults. I also include codes 466 (acute bronchitis and bronchiolitis) and 490 (bronchitis, not otherwise specified) for young children. The relevant cardiovascular diagnosis codes considered are 393-396 (heart disease), 401-405 (hypertension), 410-414 (ischemic heart disease), 427 (cardiac dysrhythmias), 428 (congestive heart failure), and 430-434 and 436-438 (cerebrovascular disease). As a check on the validity of my results, I additionally explore the effect of quarterly temperature variation on admission rates for a placebo diagnosis (i.e., one that is unlikely to be related to ambient weather): bone fractures. I identify bone fractures as those with ICD-9 prefixes of 800-829. Each hospital record contains codes for the diagnosis-upon-admission and a principal diagnosis (which is coded upon completion of the visit), which may differ from each other, as well as up to 20 secondary diagnoses. For this analysis, I consider a diagnosis to be temperature-related (or placebo-related) if any one of the possible 22 diagnosis codes is one of those listed in the relevant categories.

[5] See, for example, Lin et al. 2009.

I emphasize that the health records provided by T-DSHS and analyzed in this chapter include only the universe of *inpatient* visits, i.e. visits that resulted in at least one over-night stay in a hospital. However, it is likely that many of the morbidity impacts of extreme temperature result in outpatient visits, the records of which are not publicly available. Therefore, the estimates produced in this analysis should be understood to represent a lower-bound of the morbidity impact of extreme temperature. Assessing outpatient impacts of extreme weather is a clear direction for future research. Furthermore, to ensure that irrelevant inpatient stays (e.g., planned surgeries or other inpatient visits unlikely to be related to random weather variation) are not considered in the analysis of this paper, I limit the data to include only the inpatient visits that originated in the ER. Restricting the sample in this way results in a final database of approximately 4.7 million patient-level records over the 10-year period. Furthermore, for the analyses that investigates income via Medicaid status, I also limit the sample to ages 65 and below, which excludes low-income seniors and people with disabilities – so-called "dual eligibles" – who are enrolled in both Medicare and Medicaid simultaneously and likely have inherently different health profiles than the general population due to their advanced age.

3.3.2 Weather Data

Information on daily mean temperatures and precipitation was obtained in a 4km-by-4km grid for the contiguous U.S. from PRISM Climate Group at Oregon State University. PRISM constructs a temperature measure for each unit of the grid by interpolating daily weather station data and accounting for relevant weather factors such as wind direction. The grid was mapped to Texas geographies using shape files containing U.S. Census ZIP code tabulation areas (ZCTAs) and county borders. According to the Texas Resiliency System, there are 10 distinct "climate regions" in the state of Texas: East, Edwards Plateau, High Plains, Low Rolling Plains, Lower Valley, North Central, South, South Central, Trans Pecos, and Upper Coast. In some econometric specifications (described in Section 3.5), I control for time-invariant characteristics of these climate regions using a fixed effects estimator.

3.3.3 Additional Data

I use population data at the ZIP code level to weight the weather data to the county-level, as well as to construct admission rates for various subpopulations of interest. These data were obtained from the ACS 2015 5-year estimates. Data on the population of Medicaid enrollees was also obtained from ACS 2015 5-year estimates at the ZIP code level for a single year of my sample (2010), and is used to calculate baseline rates of admission for Medicaid patients for the temperature-related and placebo diagnosis categories listed above. Medicaid enrollee population data for earlier years will potentially be available in the future and incorporated into this analysis.

3.4 Summary Statistics

The primary unit of analysis in this study is county-quarter. There are 254 counties in Texas, yielding 10,160 county-quarter observations over the sample period 2001-2010. When records are broken out by the four primary age categores that are the focus of this paper, this increases to 40,640 observations at the county-quarter-age group level.

Desipte being an on-average hotter state, Texas exhibits a fair degree of climate variation, from arid in the west to humid in the east. Figure 3.1 summarizes the distribution of mean temperatures over the study period across each of the climate regions and seven temperature bins. Summary statistics on temperatures for each climate region have been weighted by the ZIP code population in 2010 from the ACS 2015 5-year estimates. Thus, the height of each bar represents the mean number of days of exposure per year for the average person Texan. The main empirical model in this chapter estimates a non-linear temperature-morbidity relationship using these seven bins. However, as is clear from the model, the marginal effect of temperature is restricted to be constant within each of the temperature bins. Table 3.1 compliments Figure 3.1 with summary statistics on the distribution of daily minimum, maximum and mean temperatues across the then climate regions of Texas over the study period. In particular, it shows significant variation in weather across regions, with the southern region experiencing the majority of high-temperature days (about 90°F degrees) and the plains regions experiencing the majority of the coldest days.

Figure 3.1: Annual Distribution of Daily Mean Temperature by Climate Region

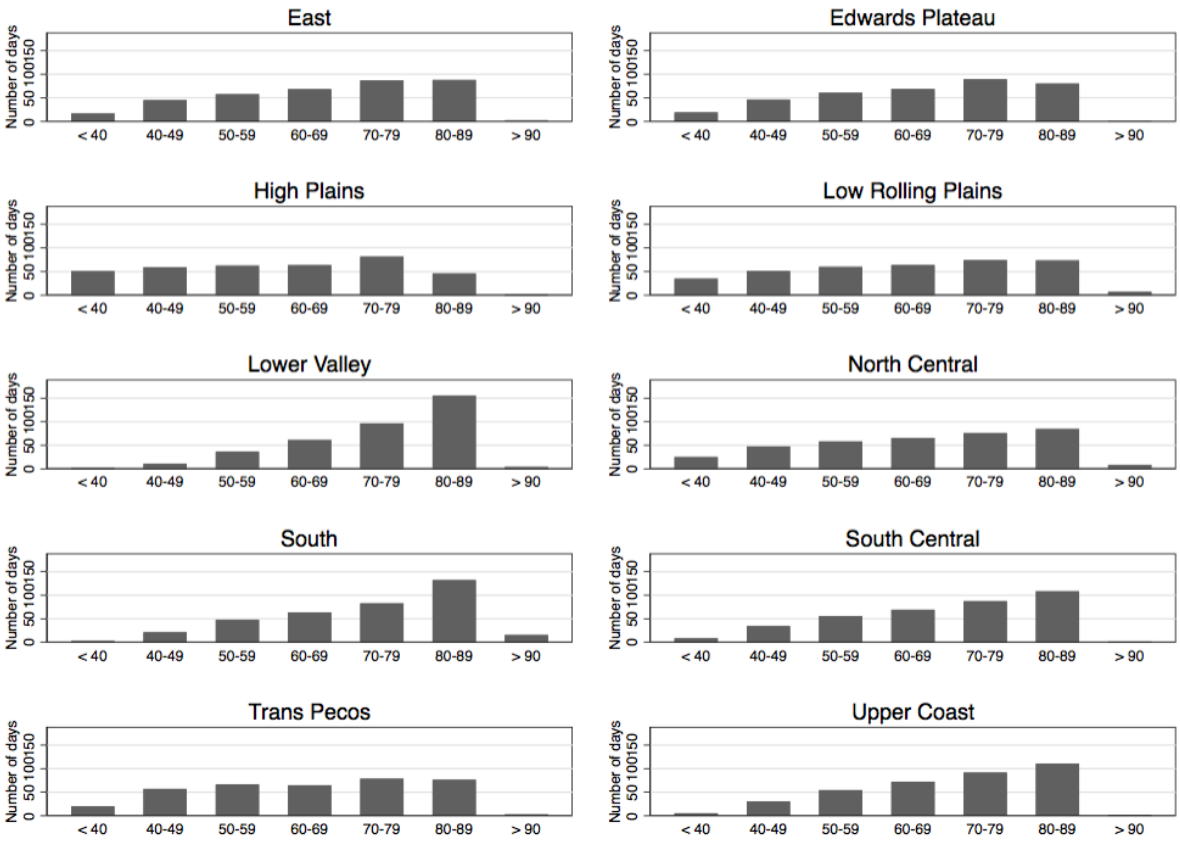

Note: Figure shows that historical average distribution of daily mean temperatures across ten temperature-day bins for each of the 10 climate regions in Texas. The bars represent the average number of days per year in each temperature category for the 1749 ZIP code-year observations over the period 1999-2010, weighted by total population in each ZIP code using the ACS 2015 5-year population estimates.

Table 3.2 presents summary statistics on the Texas patient population considered in this analysis. Note that this table contains statistics at the hospital visit-level rather than patient-level, since it is not possible to track the same individuals over time in the publicly accessibly Texas hospital records. Thus, any individuals that were admitted to a Texas hospital through the ER more than once during 2001-2010 necessarily are counted in these summary statistics equal to their total number of visits. The table reprts characteristics for three different age groups: all ages, younger than 65, and younger than 18.

The aggregate statistics reported in Table 3.2 represent total number of visits for each age-insurance category in 2010, divided by the total population in each of the categories. Data for the year 2010 is used since that is currently the single year of data with Medicaid enrollment data available at the county-level. When additional years of Medicaid enrollment data are provided by T-DSHS, this table will be updated to reflect the additional years. As is clear from the table, visit rates are higher for the Medicaid population across all categories of diagnoses, with bone fractures having the most-similar rates across the insurance types.

Table 3.1: Daily Temperatures Across 10 Texas Climate Regions, 2001-2010

	Average daily temperature			Days with mean > 90 F		Days with mean < 40 F	
Climate Region	Mean	Min	Max	Number	Mean	Number	Mean
East	66.7	55.8	77.6	1.6	90.7	17.6	35.8
Edwards Plateau	66	53.5	78.6	.7	90.6	19.1	35.1
High Plains	60.1	46.2	74	.5	90.9	51.1	32.6
Low Rolling Plains	64	51.4	76.6	6.3	91.1	35.1	33.4
Lower Valley	74.9	64.8	85	3.4	90.6	1.5	37.6
North Central	65.9	55.2	76.7	7.7	91.1	25.1	34.4
South	73.1	61.3	84.8	14.7	90.9	2.9	37
South Central	69.2	58.3	80.2	1	90.6	8.7	36.2
Trans Pecos	64.8	51.1	78.5	1.3	90.8	20	35.7
Upper Coast	69.9	60.5	79.3	.7	90.6	5.3	36.9

Notes: All statistics are weighted by ZIP code population using the ACS 2015 5-year estimates.

This may reflect the fact that Medicaid patients tend to over-utilize the ED compared to the rest of the population, as has been found in previous related work (e.g., the Oregon Health Insurance Experiment).

As described above, the Medicaid population differs significantly along key dimensions from the rest of the population. Due to the large number of children enrolled in the program, Medicaid patients are on average younger than non-Medicaid patients, even within age narrow age groups (e.g., 17 years and younger). Within children, the average age of Medicaid recipients is 5.2 years old, and for other types of insurance the average patient age is 7.4 years. Unsurprisingly, given the socioeconomic and demographic characteristics across the U.S., Medicaid patients are less often white and more often of Hispanic origin than non-Medicaid patients, across each of the three age groups considered. Interestingly, Medicaid and non-Medicaid patients reside in ZIP codes with very similar average per capita incomes–a pattern which holds across all age categories. This provides some assurance that Medicaid and non-Medicaid patients are mixed throughout similarly wealthy ZIP codes in Texas and exposed to similar climate conditions.

Risk and severity scores are coded indicators from 1 (minor) to 4 (extreme) that indicate risk of mortality and severity of illness, respectively. These scores are assigned by T-DSHS after reviewing each medical record. Risk of mortality indicates the likelihood of dying, while severity of illness indicates the extent of "physiologic decompensation". In most cases, Medicaid patients appear to be in slightly worse condition on both of these dimensions, although it is not totally clear why. It could be that, due to relatively lower reimbursement rates for Medicaid, providers are slightly less willing to admit Medicaid patients, so that the Medicaid patients who are ultimately admitted to a hospital are on average slightly sicker. Average length of stays tell a similar story: Medicaid patients are kept in the hospital slightly longer on average once they are admitted. However, I note that these differences are relatively small.

In accordance with lower reimbursement rates for Medicaid patients, total claim charges are \$500-\$1,000 lower for Medicaid patients when considering the all ages and < 65 age groups, and they are less likely to die in the hospital or be discharged to hospice care, although the rates are fairly similar. The opposite is true when looking only at children. Children with Medicaid have total claim charges that are on average over \$4,000 more than non-Medicaid claims, and they were 0.05% more likely to expire or enter hospice care than non-Medicaid children. These statistics provide further evidence that poor children are of particular interest in the temperature-health dimension since they seem to be slightly less healthy on average across all types of conditions.

Table 3.2: Characteristics of Texas ED Admit Population, 2001-2010

	All Ages		Age < 65		Age < 18	
	Non-Medicaid	Medicaid	Non-Medicaid	Medicaid	Non-Medicaid	Medicaid
Patient-level Statistics:						
Age	56.7	43.9	41.0	29.3	7.4	5.2
Male (%)	45.1%	33.0%	47.2%	35.4%	54.2%	52.9%
White (%)	69.5%	53.6%	63.8%	53.7%	62.9%	53.3%
Hispanic (%)	22.4%	36.8%	27.7%	39.4%	35.4%	47.1%
2010 ZIP Per Capita Income (\$)	27,008	27,004	26,959	26,948	26,256	26,054
Risk Score	1.75	1.71	1.41	1.44	1.12	1.17
Severity Score	2.10	2.13	1.87	1.95	1.52	1.65
Length of Stay (Days)	4.83	5.15	4.31	4.81	3.28	3.97
Total Charges (\$)	29,139	28,057	26,600	26,054	17,469	21,606
Expired or Hospice Discharge (%)	4.35%	3.78%	1.92%	1.85%	0.45%	0.50%
Aggregate Statistics:						
2010 Visit Rate: Any (%)	1.76%	2.83%	1.09%	2.22%	0.50%	0.99%
2010 Visit Rate: Respiratory (%)	0.37%	0.69%	0.17%	0.50%	0.13%	0.31%
2010 Visit Rate: Cardiovascular (%)	1.18%	1.42%	0.54%	0.74%	0.01%	0.04%
2010 Visit Rate: Bone Fractures (%)	0.04%	0.04%	0.02%	0.03%	0.01%	0.01%
Total Number of Visits (2001-2010)	6,982,530	2,320,816	3,849,917	1,593,200	359,932	467,146

Notes: All patient-level statistics are weighted by total ZIP code population using the ACS 2015 5-year estimates. The aggregate visit statistics are calculated by summing the total count of visits for the relevant diagnoses and dividing by the relevant population. Visit rates are calculated for 2010 only due to the lack of disaggregated Medicaid enrollment data for earlier years. Risk Score indicates a patient's likelihood of dying. Severity score indicates the extent of a patient's physiologic decompensation. Expired or hospise discharge indicates patients that died in the hospital or were discharged to hospice care.

3.5 Empirical Strategy

In this section, I describe the primary regression model used to estimate the relationship between ambient temperature and hospital admissions. This baseline model (not accounting for differences in insurance coverage) is identified by plausibly exogenous variation in the distribution of mean daily temperatures across counties and calendar quarters in the state of Texas. Utilizing a similar approach to that of Barreca et al. 2016, who allow for a non-linear relationship between temperature and morbidity across the different temperature bins, I estimate variants of the following model:

$$log(C_{acqy}) = \sum_j \beta_j TMEAN_{cqyj} + \gamma PRECIP_{cqy} + \delta_c + \alpha_q + \rho_y + \theta_{rq} + \phi_a + \epsilon_{acqy}, \quad (3.1)$$

where C_{acqy} is the count of hospital admissions for subgroup a in county c, quarter q and year y. County-level and region-by-quarter fixed effects are included to absorb differences

in admissions that vary across geographies (where region represents one of the ten climate regions in Texas as described above) and calendar seasons. The year indicators are included to account for idiosyncratic changes in admission rates that are common across counties in Texas, such as the rise in Medicaid enrollment during the 2009 recession. ϕ_a is a group-level fixed effect that absorbs permanent differences in admission rates that vary based on patient characteristics, such as age group.

The variables of primary interest in this study are the measures of temperature $TMEAN_{cqyj}$, which are constructed flexibly and non-parametrically and equal to the total number of days during quarter q of year y in county c in which the mean daily temperature was in the jth of 7 temperature bins in degrees Fahrenheit ($< 40°$, 40°-49°, 50°-59°, 60°-69°, 70°-79°, 80°-89°, and $> 90°$). The low-temperature bins have been collapsed to a single $< 40°$ bin due to the relative infrequency of very low mean-temperature days in Texas during 2001-2010. I omit the 60°-69°F bin as the reference category, so the coefficients on the other bins are interpreted as the morbidity impact of an additional day in the respective bin relative to an additional day in the 60°-69°F bin. The main functional form restriction implied by this parameterization is that the effect of temperature is constant within each of the bins, which previous research has shown to be a reasonable assumption when the bins are fairly narrow, as they are here. Finally, $PRECIP_{cqy}$ is equal to total quarterly precipitation in county c and year y and is included in the model based on the well-established association between ambient temperature and rainfall.

In a secondary analysis, as discussed, I attempt to explore income-related heterogeneity in the relationship between temperature and health. To do this, I augment equation (3.1) by including interactions between each of the temperature bins and Medicaid status. Specifically, I estimate

$$\begin{aligned} log(C_{acqy}) = &\sum_j \beta_j TMEAN_{cqyj} + \sum_j \pi_j TMEAN_{cqyj} \times MCAID_{acqy} \\ &+ \gamma PRECIP_{cqy} + \delta_c + \alpha_q + \rho_y + \theta_{rq} + \phi_a + \epsilon_{acqy}, \end{aligned} \tag{3.2}$$

where $MCAID_{acqy}$ equals 1 for patients with discharge records indicating a payment source of Medicaid. This specification assesses whether the effect of an additional day in a given bin is different for Medicaid (i.e., low income) patients than for non-Medicaid patients. In some additional specifications described below, I replace $MCAID_{acqy}$ with county-level per capita income or the county-level poverty rate (results not reported given their similarity to the income-based results). I also conduct an analysis in which Medicaid status is instrumented for using county-level per capita income in a two-stage least squares procedure.

3.6 Results

3.6.1 Temperature-Related Admissions for Overall Population

The regression results presented in Table 3.3 replicate findings from the epidemiology literature on temperature-related morbidity and also serve as a check on the validity of the Texas hospital data. Previous studies on the temperature-mortality relationship have tended to recover a "U-shaped" response function across the binned temperature distribution. That

is, they find an increasingly severe health response as the distribution of mean temperatures shifts farther above and below a moderate reference category (which in the present study is the 60°-69°F bin). The outcome variable for all columns of Table 3.3 is the log of the count of admissions for temperature-related diagnoses (i.e., the respiratory and cardiovascular diagnoses described in Section 3.3) in a given county and quarter for the overall universe of ER admissions resulting in an inpatient stay in Texas from 2001-2010. In contrast to previous work exploring the effect on temperature on various outcomes, in the present study the lowest temperature bins (i.e., $< 10°, 10°$-$19°, 20°$-$29°$ and $30°$-$39°$) are collapsed to a single $< 40°$ degree bin due to the relative infrequency of days in Texas where average temperature falls below 40°F.

Column (1) of Table 3.3 reports the baseline OLS estimates where the logged admission count is regressed on the seven temperature bins and, due to the close known association between temperature and rainfall, a control for total precipitation (in inches) in the respective county-quarter, as well as controls for the county-level logged poverty rate and per capita annual income. Columns (2) through (5) report the results obtained from fixed effects estimators controlling for various time-invariant trends and common economic shocks across space in Texas. Column (5) reports my preferred estimates from a model that includes fixed effects for county, year, region-by-quarter, and age group.[6]

The results reported in column (5) suggest that relative to a day with a mean temperature between 60°-69°F, days in the 70°-79°F, 80°-89°F and 90°+F bins increase the overall quarterly admission rate for temperature-related diagnoses by 0.1, 0.14 and 0.29 percentage points, respectively. Similarly, a day with a mean temperature of less than $< 40°$F increases the relevant admission rate by 0.15 percentage points. These estimates of the morbidity impacts of extreme temperature for the general population are very much in line with the existing literature on mortality impacts of temperature. For example, Deschenes and Greenstone 2011 find that relative to a day in the 50°-60°F bin, the annual mortality rate increases by 0.11 percentage points with an additional day of mean temperature $> 90°$F. Note, however, that this estimate from Deschenes and Greenstone represents the contemporaneous impact of a single additional day in the $> 90°$F bin, whereas their headline results predict that "by the end of the century climate change will lead to increases of 3 percent in the age-adjusted mortality rate." The mortality (and morbidity) effects of climate change over that time horizon will clearly be much larger than the impact of a single extreme temperature day, and I note here that climate change is not the primary focus of this paper. It is nevertheless encouraging that the single day predicted morbidity impacts for the overall population are very much in line with Deschenes and Greenstone's single day mortality impact estimates.

[6]The four age groups included in this model are the 0-17 years, 18-44 years, 45-64 years, and over 65 years.

Table 3.3: OLS Estimates — Ambient Temperature and ER Visits (Overall Population)

Temperature Bin (°F)	(1)	(2)	(3)	(4)	(5)
$< 40°$	0.0023*	0.0014***	0.0012***	0.0015***	0.0015**
	(1.93)	(3.61)	(2.72)	(2.68)	(2.24)
40°-49°	−0.0009	0.0005	0.0003	0.0006	0.0001
	(−1.00)	(1.50)	(0.69)	(1.17)	(0.26)
50°-59°	−0.0001	0.0008***	0.0008**	0.0010***	0.0013***
	(−0.10)	(2.67)	(2.09)	(2.72)	(3.37)
70°-79°	0.0000	0.0007***	0.0005**	0.0003	0.0010***
	(0.12)	(3.33)	(2.40)	(1.46)	(3.41)
80°-89°	−0.0004	0.0003**	0.0002	0.0001	0.0014***
	(−1.47)	(2.13)	(0.49)	(0.18)	(3.01)
$> 90°$	0.0011	0.0006	0.0004	0.0002	0.0029*
	(0.64)	(1.37)	(0.55)	(0.19)	(1.90)
Included controls:					
County FE		✓	✓	✓	✓
Year FE		✓	✓	✓	✓
Quarter-of-year FE			✓		
Region-by-quarter-of-year FE				✓	✓
Age group FE					✓
Observations	10,160	10,160	10,160	10,160	40,640

Note: t-statistics in parentheses. *p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the county-quarter (columns (1)-(4)) or county-quarter-age group (column (5)). The outcome variable in all regressions is the log of the count of inpatient admissions that originated in the Emergency Department and contain a respiratory and/or cardiovascular-related diagnosis. Total precipitation, per capita income (in $1,000s) and the logged poverty rate are also controlled for in each regression. Standard errors are clustered by county.

3.6.2 Temperature-Related Admissions Conditional on Medicaid Status

Table 3.4 reports the same information as Table 3.3 but includes interactions of each of the bins (as well as the omitted category bin) with an indicator variable denoting Medicaid status on the medical claim. Medicaid status is identified if either one of the two "payer source" variables indicate a payment made by Medicaid.[7] Each of the four columns report results from a fixed effects regression which includes the full set of fixed effects reported in column (5) of Table 3.3 (for county, year, region-by-quarter, and age group). The only difference between each of the four columns in Table 3.4 is the age group on which the regression is run, denoted in each of the column headings. As discussed above, there is a small population of elderly low-income patients who are insured by both Medicare and Medicaid ("dual eligibles"), although this is not very common in the data. However, since patients older than 65 are known to have different mortality and morbidity profiles than younger populations (e.g., Card, Dobkin, and Maestas 2008 find that both hospital admission rates as well as overall health status may improve with Medicare coverage), claims for visits by patients older than 65 years are excluded in all regression reported in Table 3.4.

The results in Table 3.4 suggest that Medicaid coverage is associated with fewer temperature-

[7]Most claims only have one payer source.

related admissions for all age groups except children (patients younger than 18 years). For children, I find that temperature-related admissions are higher among Medicaid recipients when there are more extreme temperature days. The opposite result is obtained for older patients, and appears to be increasing with age. That is, as Medicaid patients get older, they are less likely on average to be admitted for temperature-responsive conditions when the temperature distribution is more extreme during a given calendar quarter. As an example, relative to non-Medicaid patients, Medicaid enrollees younger than 18 see an increase in temperature-related admissions due to an additional 80°-89°F day during the calendar quarter (relative to a 60°-69°F day) of 0.07 percentage points. In contrast, Medicaid enrollees between ages 18-44 are relatively less likely to be admitted than non-enrollees by 0.20 percentage points for the same temperature shock. This admissions differential is even more extreme for patients aged 45-64, with Medicaid status associated with 0.32 percent fewer admissions in relative terms.

These results suggest that for children, the "Medicaid effect" (described previously as the empirical observation that Medicaid enrollees over-utilize emergency departments for care relative to non-Medicaid patients) outweighs the "income effect" (where higher people–those whose means exceed the eligibility limits for publicly-funded Medicaid insurance–demand relatively more health care on average). The opposite appears to be true when considering adults separately from children. These results and possible explanations are discussed in more detail in Section 3.7. I note again that these Medicaid-based estimates are suggestive rather than causal.

Table 3.4: OLS Estimates — ER Visits (Conditional on Medicaid Status)

	(1) 0-64	(2) 0-17	(3) 18-44	(4) 45-64
$< 40°$	0.0015**	0.0027**	0.0009	0.0019**
	(2.19)	(2.00)	(1.19)	(2.53)
40°-49°	0.0008	-0.0001	0.0011	0.0009
	(1.14)	(-0.06)	(1.64)	(1.29)
50°-59°	0.0009*	0.0005	0.0006	0.0013***
	(1.88)	(0.51)	(1.09)	(2.73)
70°-79°	0.0013***	0.0012*	0.0004	0.0017***
	(5.67)	(1.66)	(1.17)	(5.56)
80°-89°	0.0015***	0.0004	0.0006	0.0020***
	(3.33)	(0.63)	(1.13)	(3.82)
$> 90°$	0.0015	-0.0020	0.0003	0.0027
	(0.87)	(-1.25)	(0.17)	(1.47)
$< 40° \times$ Medicaid	-0.0022	0.0000	-0.0019*	-0.0037***
	(-1.63)	(0.01)	(-1.66)	(-2.79)
40°-49° × Medicaid	-0.0048***	-0.0005	-0.0041***	-0.0062***
	(-3.85)	(-0.27)	(-4.21)	(-4.88)
50°-59° × Medicaid	-0.0007	0.0039***	-0.0013	-0.0019*
	(-0.70)	(3.40)	(-1.54)	(-1.70)
60°-69° × Medicaid	-0.0024***	0.0014**	-0.0020***	-0.0033***
	(-6.52)	(2.05)	(-5.64)	(-8.85)
70°-79° × Medicaid	-0.0034***	0.0003	-0.0025***	-0.0044***
	(-19.57)	(0.87)	(-11.36)	(-19.97)
80°-89° × Medicaid	-0.0026***	0.0007**	-0.0020***	-0.0032***
	(-7.77)	(2.48)	(-5.88)	(-8.81)
$> 90° \times$ Medicaid	-0.0003	0.0031	-0.0006	-0.0009
	(-0.16)	(1.34)	(-0.32)	(-0.61)
Observations	20,320	20,320	20,320	20,320

Note: t-statistics in parentheses. *p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the county-quarter-Medicaid status (a binary categorical variable) for the respective age group. The outcome variable in all regressions is the log of the count of inpatient admissions that originated in the Emergency Department and contain a respiratory and/or cardiovascular-related diagnosis. Total precipitation, per capita income (in $1,000s) and the logged poverty rate are also controlled for in each regression. Standard errors are clustered by county.

3.6.3 Temperature-Related Admissions Conditional on County-Level Per Capita Income

As discussed previously, health care is a superior good, and previous research has shown large income elasticities in the demand for health care, which could partially explain the pattern observed in Table 3.4, particularly if higher income adults (i.e., those less likely to be insured by Medicaid) are more likely to need (or at least, to seek) health care during periods of extreme temperatures. Table 3.5 reports comparable results to Table 3.4 using an alternative measure of income than Medicaid status: county-level per capita income (reported in $1,000s). These results show that income is positively associated with hospital admissions for children, and negatively associated with admissions for adults, which is the opposite result from the effect obtained in the previous table if one considers Medicaid as a pure proxy for income. These results together suggest that Medicaid recipiency and income work in opposite directions when considering effects for children and adults separately. For children, having Medicaid (as opposed to any other form of insurance) makes them more likely to seek medical care on very hot or very cold days. Yet, in a contrasting result, higher county-level income (which is certainly negatively correlated with likelihood of Medicaid status among county residents) is also associated with increased hospital utilization by children during the same periods. The opposite holds for adults. In any case, it appears that for children, the county-level income effect (an approximate increased in admissions of 0.3 percentage points for each of the temperature bins) outweighs the Medicaid effect (approximately 0.01 percentage points for each of the bins).

Given previous research showing that Medicaid status is likely correlated with unobservables that influence both the level of demand for emergency care as well as the likelihood of hospital admission (assuming that providers have a degree of discretion in which patients are ultimately admitted), it is likely that the estimates in Table 3.4 suffer from omitted variables bias. Table 3.6 reports results of a two stage least squares regression in which Medicaid is instrumented in a first stage model with county-level income. (Table 3.5 thus serves as a sort of reduced form depiction of this relationship). The IV results essentially pick up no effect of instrumented Medicaid status on the log of total admissions for the relevant diagnoses. The first-stage results (not reported) indicate that county-level income is (perhaps not surprisingly) a fairly weak instrument for Medicaid status of individual residents. Indeed, the F-statistic in the first stage is less than 1 for most of the age groups, which clearly underscores a weak instruments problem. In future iteratures of this work, it could be enlightening to run a similar two stage least squares analysis instead relying on a more granular income measure, perhaps at the ZIP code, Census Tract, or Census Block level. Income at one of these lower levels of granularity would likely be a better predictor of an individual's Medicaid status, given the degree of income variation one might expect to exist in geographic areas as large as a county.

Table 3.5: OLS Estimates — ER Visits (Conditional on Per Capita Income)

	(1) 0-64	(2) 0-17	(3) 18-44	(4) 45-64
$< 40°$	-0.0021	-0.0074***	-0.0029	-0.0008
	(-1.41)	(-2.70)	(-1.63)	(-0.49)
40°-49°	-0.0012	0.0001	0.0020	-0.0026**
	(-1.15)	(0.02)	(1.19)	(-2.15)
50°-59°	-0.0025**	-0.0049*	-0.0040**	-0.0011
	(-2.12)	(-1.75)	(-2.38)	(-0.95)
70°-79°	-0.0026***	-0.0045**	-0.0018	-0.0023**
	(-2.86)	(-2.58)	(-1.33)	(-2.32)
80°-89°	-0.0001	-0.0012	-0.0004	0.0004
	(-0.06)	(-0.78)	(-0.30)	(0.41)
$> 90°$	0.0035	0.0042	0.0036	0.0031
	(1.43)	(1.32)	(1.35)	(1.19)
$< 40°$ * Inc.	-0.0004	0.0034***	-0.0007	-0.0011**
	(-1.01)	(4.57)	(-1.57)	(-2.10)
40°-49° × Income	-0.0005	0.0031***	-0.0009*	-0.0011**
	(-1.24)	(4.41)	(-1.93)	(-2.15)
50°-59° × Income	-0.0004	0.0033***	-0.0007	-0.0011**
	(-1.03)	(4.79)	(-1.52)	(-2.15)
60°-69° × Income	-0.0005	0.0031***	-0.0009*	-0.0012**
	(-1.30)	(4.41)	(-1.88)	(-2.25)
70°-79° × Income	-0.0004	0.0033***	-0.0008*	-0.0011**
	(-1.06)	(4.64)	(-1.70)	(-2.10)
80°-89° × Income	-0.0005	0.0032***	-0.0008*	-0.0012**
	(-1.23)	(4.54)	(-1.81)	(-2.22)
$> 90°$ × Income	-0.0006	0.0030***	-0.0009*	-0.0012**
	(-1.37)	(4.33)	(-1.86)	(-2.24)
Observations	10,160	10,160	10,160	10,160

Note: t-statistics in parentheses. *p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the county-quarter-Medicaid status (a binary categorical variable) for the respective age group. The outcome variable in all regressions is the log of the count of inpatient admissions that originated in the Emergency Department and contain a respiratory and/or cardiovascular-related diagnosis. Total precipitation and the logged poverty rate are also controlled for in each regression. Standard errors are clustered by county.

Table 3.6: IV Estimates — ER Visits (Instrumenting for Medicaid with Per Capita Income)

	(1) 0-64	(2) 0-17	(3) 18-44	(4) 45-64
$< 40°$	0.0090	0.0398	0.0068	0.0108
	(1.09)	(0.82)	(0.29)	(0.45)
$40° - 49°$	0.0064	-0.0360	0.0130	0.0169
	(0.48)	(-0.63)	(0.24)	(0.50)
$50° - 59°$	-0.0051	0.1405	-0.0386	-0.0162
	(-0.24)	(0.96)	(-0.37)	(-0.61)
$70° - 79°$	0.0067	0.0211	0.0037	0.0032
	(1.15)	(0.71)	(0.16)	(0.40)
$80° - 89°$	-0.0075	0.0894	-0.0305	-0.0088
	(-0.54)	(0.83)	(-0.48)	(-0.76)
$> 90°$	-0.0166	0.0872	-0.0636	-0.0177
	(-0.63)	(0.68)	(-0.55)	(-0.68)
$< 40° \times \widehat{\text{Medicaid}}$	-0.0075	-0.0594	0.0168	-0.0168
	(-0.33)	(-0.59)	(0.18)	(-0.25)
$40° - 49° \times \widehat{\text{Medicaid}}$	-0.0029	0.0677	0.0121	-0.0398
	(-0.11)	(0.64)	(0.11)	(-0.43)
$50° - 59° \times \widehat{\text{Medicaid}}$	0.0323	-0.2339	0.1490	0.0704
	(0.51)	(-1.03)	(0.42)	(0.76)
$60° - 69° \times \widehat{\text{Medicaid}}$	0.0141	0.0143	0.0322	0.0128
	(1.15)	(0.23)	(0.56)	(0.59)
$70° - 79° \times \widehat{\text{Medicaid}}$	-0.0049	-0.0157	0.0145	0.0032
	(-0.43)	(-0.44)	(0.28)	(0.23)
$80° - 89° \times \widehat{\text{Medicaid}}$	0.0327	-0.1141	0.1050	0.0367
	(0.96)	(-0.82)	(0.55)	(1.24)
$> 90° \times \widehat{\text{Medicaid}}$	0.0499	-0.0891	0.1678	0.0560
	(0.93)	(-0.61)	(0.61)	(0.94)
Observations	20,320	20,320	20,320	20,320

Note: t-statistics in parentheses. *p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the county-quarter-Medicaid status for the respective age group. Medicaid status is instrumented using county-level per capita income. The outcome variable in all regressions is the log of the count of inpatient admissions that originated in the Emergency Department and contain a respiratory and/or cardiovascular-related diagnosis. Total precipitation and the logged poverty rate are also controlled for in each regression. Standard errors are clustered by county.

I conduct several supplemental robustness checks, which are reported in Appendix Section C.1, in order to gauge the stability of these estimates. First, I conduct a placebo test to ensure that the estimated morbidity impacts are indeed a response to short-run temperature

variation, rather than spurrious correlation or due to some other unobservable factor. This type of placebo test is common practice in the literature on environmental health. For this analysis, I run a placebo check using a medical diagnosis that theoretically should not respond to ambient temperature: bone fractures. For this test, I limit the sample to only children because they seem to be the age group least likely to suffer from insurance-related selection issues (i.e., if the child genuinely needs medical care, parents/guardians will likely seek it). These estimates confirm that hospitalization rates for bone fractures are not empirically associated with extreme temperature. The two columns in Table C.1 report results from regressions which include the full set of fixed effects used in column (4) of Table 3.3 (for county, year, and region-by-quarter). The lack of statistical significance of essentially all the estimated coefficients in Table C.1 suggests that the morbidity impacts estimated in each of these models are likely a geniune response to climate variation.

Second, to address potential selection effects by income and Medicaid status, I narrow the focus to very severe illnesses that theoretically would require hospitalization regardless of financial or other demand-side constraints. These results are reported in Table C.2. These results suggest that admissions are relatively lower for Medicaid patients for severe diagnoses during periods of extreme temperature, particularly when focusing on older patients. The impact is largest for patients aged 45-64, who experience approximately 0.2 percent fewer admissions for severe conditions. Based upon the (potentially quite strong) assumption that the demand for health care should be universal for these illnesses, this provides at least weak evidence that Medicaid-insured adults fare better against climate shocks than patients with other forms of insurance. Of course, it might be the case that Medicaid patients suffering severe illness in response to extreme temperature more frequently die prior to reaching the ER, or conversely, are treated fairly quickly in the ER without being admitted to the hospital. In either case, these patients would not appear in the hospital inpatient sample that is used in this analyiss. In future work, it could be worthwhile to examine whether there is a corresponding *increase* in mortality and/or outpatient (non-admitted) ER visits among Medicaid-insured adults in order to rule out these possibilities.

Finally, I also run the main Medicaid analysis at the ZIP code rather county level to see whether improving the geographic resolution of the data has any impact on the estimated effect sizes.[8] Qualitatively, the results are unchanged from the county-level analysis.

3.7 Discussion and Conclusion

Using Texas as a case study, the results of this analysis confirm that increased morbidity, like mortality, is a health consequence of extreme ambient temperature. Additionally, while lacking strong causal foundations, the supplemental Medicaid results presented in this chapter suggest that both public health insurance and county-level income drive statistically significant differences in climate-related morbidity, although relative impacts differ across age groups. Relatively higher admissions among Medicaid-insured children in response to extreme temperature could indicate a decreased ability in this population to endure extreme weather conditions – perhaps through lower baseline health stock or more hours spent out-

[8]It was not possible to run the income-based and poverty-based analyses at the ZIP level due to lack of publicly available data for these variables at the time of analysis.

doors on average – as compared to children with other forms of insurance (and likely higher household income). Indeed, as summarized in Section 3.2, the body's thermo-regulatory system is less developed in children, indicating that they are likely more vulnerable to their surrounding environment than adults. If children in Medicaid families are more likely to be exposed to extreme temperature, this could be an explanation for the results obtained in this analysis. On the other hand, relevant admissions for children in higher-income counties are relatively higher than in lower-income counties, perhaps due to a higher medical spending elasticity in this population, or perhaps a lower insurance penetration among the poor and thus lower overall health care utilization. Although the underlying reasons for these differences have not been identified, and the results come with several important caveats, these findings point to the likelihood that both income and public health insurance can moderate the health impacts of extreme climate.

Interestingly, while acknowledging the clear limitations of the current research design, it appears that the income and Medicaid effects work in opposite directions, with the income effect relatively larger in magnitude for children. If hospital utilization is considered a reasonable proxy for underlying health status, then the relatively larger health damages among children in wealthier counties estimated in this chapter run counter to the mortality estimates in Carleton et al. 2020, who find that income can mitigate the damaging health impacts of extreme climate. More work should be conducted in this area to disentangle these two competing explanations, including leveraging quasi-random variation in insurance coverage (perhaps by instead focusing on a state which did adopt the ACA Medicaid expansion and using an exposure-based instrument, as in Chapter 1 of this dissertation). The selection bias that almost certainly underpins the current estimates likely derives from two separate channels: supply-side constraints (e.g., limits on hospital capacity, lower prevalence of providers that accept Medicaid insurance, etc.), and demand-side factors (e.g., differential demand for health care across both insurance plans and the income distribution in general). Exploring these possibilities further would require data on hospital characteristics, such as location (e.g., ZIP code) and capacity (number of beds, physician availability, etc.). These data exist and have been requested from the Texas hospital association through an open records request, but the request has not yet been fulfilled. In future work, it may be interesting to explore distance-based heterogeneity in hospital admissions (i.e., comparing patients who reside in ZIP codes within 5 miles from the hospital, 5-10 miles, and so on) to assess whether access may play a role in the estimated impacts. It could also be useful to control for county-level hospital capacity, which could shed light on the potential issue of capacity constraints.

This study provides suggestive evidence on the temperature-health relationship for potentially vulnerable subpopulations: those that qualify for Medicaid and those that are low-income in general. To my knowledge, it is the first study of its kind to exploit individual-level health care claims data to investigate income-based (although non-causal) heterogeneity in the temperature-health relationship. Preferred estimates suggest that temperature-related admission rates are relatively higher among children insured by Medicaid than children with all other forms of insurance. Conversely, admissions are relatively lower for Medicaid-insured adults, a difference which might be explained by an improved ability to endure extreme temperature that arises from long-term participation in the Medicaid program (although, with these data, it is currently only possible to determine Medicaid status at the point in time of

a patient's hospital visit). However, if improved climate protection is determined to indeed be partially attributable to publicly funded health insurance, then the health care-related policy recommendations are clear: incorporate these additional insurance benefits into future benefit-cost analyses around public insurance expansions. Additional empirical work based on exogenous variation in health insurance coverage is necessary to determine with certainty whether improved resiliency to climate variation is one unintended benefit of the Medicaid program, and this paper is a first step in exploring that possibility.

References

Anderson, M., C. Dobkin, and T. Gross. 2012. "The Effect of Health Insurance Coverage on the Use of Medical Services". *American Economic Journal: Economic Policy* 4 (1): 1–27.

Aron-Dine, A., L. Einav, and A. Finkelstein. 2012. "The RAND Health Insurance Experiment, Three Decades Later". NBER Working Paper 18642.

Aufhammer, M, V Ramanathan, and J Vincent. 2012. "Climate change, the monsoon, and rice yield in India". *Climate Change* 111:411–424.

Barreca, A., et al. 2016. "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century". *Journal of Political Economy* 124 (1): 105–159.

Becker, Gary. 1965. "A theory of the allocation of time". *Economic Journal* 75 (299): 493–517.

Bishop, K, J Ketcham, and N Kuminoff. 2018. "Hazed and Confused: Air Pollution, Dementia, and Financial Decision Making". *NBER Working Paper No. 24970*.

Burke, M, S Hsiang, and E Miguel. 2015. "Global non-linear effect of temperature on economic production". *Nature* 527 (7577): 235–239.

Burke, M., et al. 2020. "The Changing Risk and Burden of Wildfire in the US". *NBER Working Paper No. 27423*.

Burke, M, et al. 2009. "Warming increases the risk of civil war in Africa". *Proceedings of the National Academy of Sciences* 106:20670–20674.

Card, D., C. Dobkin, and M. Maestas. 2008. "The Impact of Nearly Universal Insurance Coverage on Health Care Utilization: Evidence from Medicare". *American Economic Review* 98 (5): 2242–2258.

Card, D., C. Dobkin, and N. Maestas. 2004. "The Impact of Nearly Universal Insurance Coverage on Health Care Utilization and Health: Evidence from Medicare". *NBER Working Paper No. 10365*.

Carleton, T., and S. Hsiang. 2016. "Social and economic impacts of climate". *Science* 353 (6304).

Carleton, T., et al. 2020. "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits". *NBER Working Paper No. 27599*.

Cascio, W. 2019. "Wildfire Smoke and Public Health - Why is the EPA Concerned? Web Summit." United States Environmental Protection Agency, Office of Research and Development.

Center on Budget and Policy Priorities (CBPP). 2020. *The Far-Reaching Benefits of the Affordable Care Act's Medicaid Expansion*.

Deryugina, T. 2017. "The Fiscal Cost of Hurricanes: Disaster Aid versus Social Insurance". *American Economic Journal: Economic Policy* 9 (3): 168–198.

Deryugina, T., et al. 2019. "The Mortality and Medical Costs of Air Pollution: Evidence from Changes in Wind Direction". *NBER Working Paper 22796*.

Deschenes, O., and M. Greenstone. 2011. "Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US". *American Economic Journal: Applied Economics* 3 (4): 152–85.

Deschenes, O., M. Greenstone, and J. Shapiro. 2017. "Defensive Investments and the Demand for Air Quality: Evidence from the NOx Budget Program". *American Economic Review* 107 (10): 2958–2989.

Donahue, J., et al. 1997. "Inhaled Steroids and the Risk of Hospitalization for Asthma". *JAMA* 277 (11): 887–891.

Finkelstein, A., et al. 2012. "The Oregon Health Insurance Experiment: Evidence from the First Year". *Quarterly Journal of Economics* 127 (3): 1057–1106.

Graff-Zivin, J, S Hsiang, and M Neidell. 2018. "Temperature and human capital in the short- and long-run". *Journal of the Association of Environmental and Resource Economists* 5 (1): 77–105.

Grossman, M. 1972. "On the concept of health capital and the demand for health". *Journal of Political Economy*: 223–255.

Hall, R, and C Jones. 2007. "The Value of Life and the Rise in Health Spending". *Quarterly Journal of Economics* 122:39–72.

Henneman, Lucas, Loretta Mickley, and Cory Zigler. 2019. "Air pollution accountability of energy transitions: the relative importance of point source emissions and wind fields in exposure changes". *Environmental Research Letters*.

Heutel, G, N Miller, and D Molitor. 2020. "Adaptation and the Mortality Effects of Temperature Across U.S. Climate Regions". *NBER Working Paper 23271*.

Holland, Stephen, et al. 2020. "Decompositions and Policy Consequences of an Extraordinary Decline in Air Pollution from Electricity Generation". *American Economic Journal: Economic Policy* 12 (4): 244–274.

Jaglom, W, et al. 2014. "Assessment of projected temperature impacts from climate change on the U.S. electric power sector using the Integrated Planning Model". *Energy Policy* 73:524–539.

Kaiser Family Foundation. 2020. *Medicaid State Fact Sheets*. Online.

Krewski, Daniel, et al. 2009. *Extended Follow-Up and Spatial Analysis of the American Cancer Society Study Linking Particulate Air Pollution and Mortality*. HEI Research Report 140. Boston, MA: Health Effects Institute.

Lepeule, Johanna, et al. 2012. "Chronic Exposure to Fine Particles and Mortality: An Extended Follow-Up of the Harvard Six Cities Study from 1974 to 2009". *Environmental Health Perspectives* 120 (7): 965–70.

Lin, S, et al. 2009. "Extreme High Temperatures and Hospital Admissions for Respiratory and Cardiovascular Diseases". *Epidemiology* 20 (5).

Merchant, R., et al. 2018. "Impact of a digital health intervention on asthma resource utilization". *World Allergy Organization Journal* 11 (28).

Miller, N., D. Molitor, and E. Zou. 2017. "Blowing Smoke: Health Impacts of Wildfire Plume Dynamics". Working Paper.

Miller, S. 2012. "The effect of insurance on emergency room visits: An analysis of the 2006 Massachusetts health reform". *Journal of Public Economics* 96:893–908.

Mullins, J., and C. White. 2020. "Can Access to Health Care Mitigate the E ects of Temperature on Mortality?" *Working Paper*.

OECD. 2016. "The Economic Consequences of Outdoor Air Pollution".

Public Policy Institute of California (PPIC). 2018. *The Affordable Care Act in California*. Online.

Ranson, M. 2014. "Crime, weather and climate change". *Journal of Environmental Economics and Management* 67 (3): 274–302.

Schlenker, W, and M Roberts. 2009. "Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change". *Proceedings of the National Academy of Sciences* 106 (37): 15594–15598.

Schlenker, W., and W. R. Walker. 2016. "Airports, Air Pollution, and Contemporaneous Health". *The Review of Economic Studies* 83 (2): 768–809.

Shapiro, Joseph, and W. Reed Walker. 2020. "Is Air Pollution Regulation Too Stringent?" Working Paper.

Sin, D., et al. 2004. "Pharmacological Management to Reduce Exacerbations in Adults With Asthma: A Systematic Review and Meta-analysis". *JAMA* 292 (3): 367–376.

Stock, J., M Yogo, and J Wright. 2002. "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments". *Journal of Business and Economic Statistics* 20:518–529.

Taubman, S, et al. 2014. "Medicaid Increases Emergency-Department Use: Evidence from Oregon's Health Insurance Experiment". *Science* 343 (6168): 263–268.

Tessum, Christopher W., Jason D. Hill, and Julian D. Marshall. 2017. "InMAP: A model for air pollution interventions". *PLOS ONE* 12 (4).

U.S. EPA. 2011a. *Regulatory Impact Analysis for the Final Mercury and Air Toxics Standards*. Technical Report. EPA-452/R-11-011. Office of Air Quality Planning and Standards.

— . 2012. *Regulatory Impact Analysis for the Final Revisions to the National Ambient Air Quality Standards for Particulate Matter*. Technical Report. EPA-452/R-12-005. Office of Air Quality Planning and Standards, Research Triangle Park, NC.

— . 2011b. *Regulatory Impact Analysis (RIA) for the final Transport Rule*. Technical Report. EPA-HQ-OAR-2009-0491. Office of Air and Radiation.

Wang, J., and A. Trivedi. 2017. "Enrollment in California's Medicaid Program After the Affordable Care Act Expansion". *American Journal of Public Health* 107 (11): 1757–1759.

Wang, T., et al. 2014. "Emergency Department Charges for Asthma-Related Outpatient Visits by Insurance Status". *Journal of Health Care for the Poor and Underserved* 25 (1): 396–405.

Wang, Y, et al. 2020. "Spatial decomposition analysis of NO_2 and $PM_{2.5}$ air pollution in the United States". *Atmospheric Environment* 241.

Williams, A. Park, et al. 2019. "Observed Impacts of Anthropogenic Climate Change on Wildfire in California". *Earth's Future* 7:892–910.

Appendix A

Supplemental Materials for Chapter 1

A.1 Data Notes

A.1.1 Identification of Selected Health Conditions

Table A.1: Diagnostic Code Prefixes Associated with Selected Medical Conditions

Diagnosis	ICD-9	ICD-10
Asthma	493	J44, J45
All Respiratory	460-466, 470-488, 490-496, 500-519	J00-J18, J20-J22, J30-J47, J60-J70, J80-J86, J90-J99
Hypertension	401, 402	I10, I11
Heart Attack	410	I21
Diabetes	250	E10, E11
Appendicitis	540	K35

Note: Table reports the prefixes in diagnostic codes used to identify health conditions under the International Classification of Diseases (ICD) coding system. Hospitals switched from ICD-9 to ICD-10 codes in the fourth quarter of 2015.

A.1.2 ZIP Codes vs. ZCTAs

Throughout the paper, I refer interchangeably to ZIP codes and ZIP Code Tabulation Areas as simply ZIPs. However, ZIP codes are linear entities, representing a collection of United States Postal Service (USPS) routes, while ZCTAs are geo-coded area representations of 5-digit ZIP Code service areas. The Census Bureau created ZCTAs to map more closely to block-level information collected in surveys and censuses. The Census Bureau assigns 5-digit ZCTAs based on the most frequently occurring USPS ZIP Code within that ZCTA. Therefore, ZCTAs may contain more than one ZIP Code. The Census Bureau does not provide an official crosswalk between ZIPs and ZCTAs, but other entities do. I map ZIP Codes to ZCTAs using the crosswalk provided here: https://www.udsmapper.org/zcta-crosswalk.cfm.

A.2 Derivation of Comparative Statics

A.2.1 Model Setup

Model summary:

- Two-period model (no time discounting)
- Maximize sum of utility in both periods subject to budget constraint
- Insurance status determined outside model and sets prices paid for healthcare (preventive and emergency)
- Pollution shock in $t = 2$, with level determined exogenously.

Variables:

- Non-healthcare consumption: x
- Preventive care (purchased in $t = 1$): q
- Emergency care (purchased in $t = 2$): e
- Pollution shock shock (experienced in $t = 2$): c
- Prices paid for preventive and emergency care: p_q, p_e

Health production function:

- $H = H(q, e, c)$
- Pollution is bad: $\partial H/\partial c < 0$
- Healthcare is good: $\partial H/\partial q > 0$, $\partial H/\partial e > 0$

A.2.2 Optimization Problem

Constrained maximization:

$\max_{x,q,e} U(x) + H(q, e, c)$ s.t. $x + p_q q + p_e e \leq Y$

First order conditions:

1. $U_x - \lambda = 0$
2. $H_q - p_q\lambda = 0$
3. $H_e - p_e\lambda = 0$
4. $Y - x - p_q q - p_e e = 0$

Total differentials of the four FOCs:

$$\begin{bmatrix} U_{xx} & 0 & 0 & -1 \\ 0 & H_{qq} & H_{qe} & -p_q \\ 0 & H_{qe} & H_{ee} & -p_e \\ 1 & p_q & p_e & 0 \end{bmatrix} \begin{bmatrix} dx \\ dq \\ de \\ d\lambda \end{bmatrix} = D \begin{bmatrix} dx \\ dq \\ de \\ d\lambda \end{bmatrix} = \begin{bmatrix} 0 \\ -H_{qc}dc + \lambda dp_q \\ -H_{ec}dc + \lambda dp_e \\ -qdp_q - edp_e \end{bmatrix} \tag{1}$$

Applying Cramer's Rule to solve for de:

$|D| = U_{xx}(p_q^2 H_{ee} + p_e^2 H_{qq} - 2p_q p_e H_{qe}) + H_{qq}H_{ee} - H_{qe}^2$

$|D| > 0$ by SOC for a maximum.

By Cramer's Rule,

$$de = \frac{\begin{vmatrix} U_{xx} & 0 & 0 & -1 \\ 0 & H_{qq} & -H_{qc}dc + \lambda dp_q & -p_q \\ 0 & H_{eq} & -H_{ec}dc + \lambda dp_e & -p_e \\ 1 & p_q & -qdp_q - edp_e & 0 \end{vmatrix}}{|D|} = \frac{C \cdot dc + Q \cdot dp_q + E \cdot dp_e}{|D|} \tag{2}$$

where

$$C = U_{xx}(p_q p_e H_{qc} - p_q^2 H_{ec}) + (H_{qe}H_{qc} - H_{qq}H_{ec})$$

$$Q = \lambda p_q p_e - p_e q H_{qq} + p_q q H_{qe} + \lambda H_{qe}$$

$$E = -\lambda p_q^2 + p_q e H_{qe} - p_e e H_{qq} + \lambda H_{qq}$$

Then, assuming only c changes,

$$\frac{de}{dc} = \frac{U_{xx}(p_q p_e H_{qc} - p_q^2 H_{ec}) + (H_{qe}H_{qc} - H_{qq}H_{ec})}{|D|} \tag{3}$$

A.3 Insurance Variation (County-Level)

Figure A.1: Adult Uninsured Rates, by County (2013)

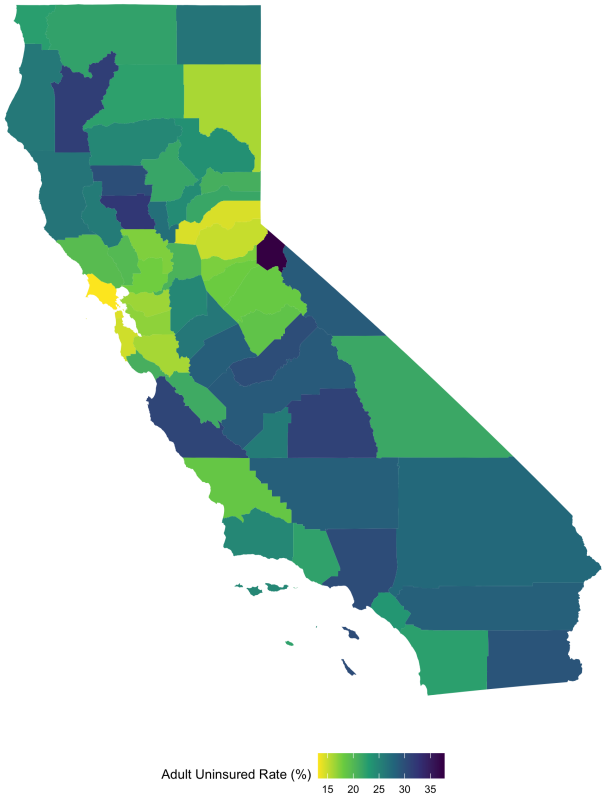

Notes: Figure shows the county-level rate of uninsured adults in 2013. Data come from the American Community Survey 5-Year estimates.

A.4 Summary Statistics (ZIP-Week)

Table A.2: Summary Statistics, 2012–2017 (ZIP-Week)

	N	Mean	p25	p75	Max
Morbidity outcomes:					
Total ER visits	560,316	142.9	10	218	3,276
ER visits: Asthma	560,316	2.3	0	3	135
ER visits: All Respiratory	560,316	14.8	1	20	675
ER visits: Hypertension	560,316	0.9	0	1	17
ER visits: Heart Attack	560,316	0.6	0	1	13
ER visits: Diabetes	560,316	1.3	0	2	23
ER visits: Appendicitis	560,316	0.4	0	1	11
ER visits: Bone Fractures	560,316	1.4	0	2	18
Air quality (Jun–Aug):					
% smoke coverage	137,436	20.8	0	31.0	100
Mean $PM_{2.5}$	11,070	9.5	6.1	11.7	195.6
Maximum $PM_{2.5}$	11,070	13.7	8.3	16.0	498.0
Air quality (Sep–Nov):					
% smoke coverage	137,436	2.2	0	0	100
Mean $PM_{2.5}$	10,994	10.1	5.2	12.5	232.1
Maximum $PM_{2.5}$	10,994	15.3	7.5	19.1	557.0
Air quality (Dec–Feb):					
% smoke coverage	137,436	0.3	0	0	69
Mean $PM_{2.5}$	11,106	9.3	4.7	11.6	481.0
Maximum $PM_{2.5}$	11,106	14.2	6.9	17.9	985.0
Air quality (Mar–May):					
% smoke coverage	137,436	3.4	0	0	100
Mean $PM_{2.5}$	11,101	7.9	5.0	10.1	52.0
Maximum $PM_{2.5}$	11,101	11.1	7.1	13.8	167.3
Population and health insurance:					
Total population	560,316	21,378	1,497	35,790	107,888
Adult population	560,316	13,497	895	22,609	67,363
% Uninsured adults, 2013 (Z)	546,960	23.8	14.2	31.6	100.0
% Medicaid coverage	551,019	28.4	13.6	39.1	100.0

Notes: Table reports unweighted statistics for the main morbidity estimation sample. Unit of observation is the ZIP-week. $PM_{2.5}$ data is only available for the ZIP codes that have official EPA monitors. Mortality data is not available at an aggregation lower than county-by-month.

A.5 IV-DID Model Assumptions

A.5.1 Parallel Trends

Table A.3: Pre-Trends Analysis — ER Visits (Selected Conditions)

	Selected Conditions				Selected Placebos		
	Asthma	All Resp.	Hypertension	Heart Attack	All Conditions	Diabetes	Appendicitis
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Smoke	0.018***	−0.0002	0.002	0.002	0.0004	−0.004	0.002
	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)	(0.002)
Medicaid	0.019*	0.031***	0.024***	0.018***	0.057***	0.023***	0.004
	(0.010)	(0.010)	(0.009)	(0.005)	(0.011)	(0.007)	(0.004)
Smoke × Medicaid × 2012	0.001	0.013***	0.003	−0.001	0.001	0.001	−0.003
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Smoke × Medicaid × 2014	0.008***	0.007***	0.001	0.001	0.005***	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Smoke × Medicaid × 2015	−0.009***	−0.006***	−0.002*	−0.001	−0.002**	0.003***	−0.0002
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Smoke × Medicaid × 2016	−0.004***	−0.005***	−0.006***	0.001	−0.004***	0.002*	−0.0004
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Smoke × Medicaid × 2017	−0.009***	−0.008***	−0.004***	−0.001	−0.002*	0.002*	−0.002*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	545,219	545,219	545,219	545,219	545,219	545,219	545,219
Adjusted R^2	0.772	0.929	0.557	0.384	0.977	0.681	0.337

Note: *p<0.1; **p<0.05; ***p<0.01. The dependent variable is the log of the count of ER visits for the stated condition in ZIP j and week t of year y. Medicaid is instrumented for using the rate of uninsured adults in the ZIP code in 2013 (as in the ER visits models in the main text), and Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across ZIP-weeks (equal to 17.2% coverage) or greater. The excluded year is 2013. All models include ZIP-week-of-year and year fixed effects and weather and socioeconomic controls. Standard errors are clustered by ZIP.

Table A.4: Pre-Trends Analysis — Mortality (Selected Conditions)

	Selected Conditions				Selected Placebos		
	Asthma	All Resp.	Hypertension	Heart Attack	All Conditions	Diabetes	Appendicitis
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Smoke	0.025	−0.007	−0.009	0.036	0.016	−0.003	0.018
	(0.043)	(0.039)	(0.041)	(0.043)	(0.033)	(0.043)	(0.018)
Medicaid	−0.0005	0.029	−0.020	0.108	0.042	−0.092	0.001
	(0.060)	(0.057)	(0.110)	(0.117)	(0.054)	(0.139)	(0.021)
Smoke × Medicaid × 2012	0.056	0.083***	0.016	0.012	0.026	0.059*	−0.004
	(0.037)	(0.030)	(0.031)	(0.039)	(0.028)	(0.032)	(0.007)
Smoke × Medicaid × 2014	0.018	0.022	0.028	−0.026	0.016	0.024	−0.004
	(0.025)	(0.026)	(0.020)	(0.018)	(0.015)	(0.016)	(0.006)
Smoke × Medicaid × 2015	0.001	0.005	0.015	−0.011	0.007	0.008	0.001
	(0.017)	(0.017)	(0.015)	(0.016)	(0.012)	(0.016)	(0.005)
Smoke × Medicaid × 2016	0.009	0.009	−0.012	−0.023	−0.001	0.007	−0.001
	(0.015)	(0.016)	(0.016)	(0.016)	(0.011)	(0.014)	(0.006)
Smoke × Medicaid × 2017	−0.010	−0.012	0.004	−0.010	−0.0002	0.011	−0.007*
	(0.013)	(0.012)	(0.017)	(0.014)	(0.012)	(0.015)	(0.004)
Observations	4,176	4,176	4,176	4,176	4,176	4,176	4,176
Adjusted R^2	0.925	0.951	0.916	0.917	0.989	0.900	0.142

Notes: *p<0.1; **p<0.05; ***p<0.01. The dependent variable is the log of the count of deaths for the stated condition in county j and week t of year y. Medicaid is instrumented for using the rate of uninsured adults in the county in 2013 (as in the mortality models in the main text), and Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across county-months (equal to 15.4% coverage) or greater. The excluded year is 2013. All models include county-month-of-year and year fixed effects and weather and socioeconomic controls. Standard errors are clustered by county.

A.5.2 IV First Stage: Insurance (County-Level)

Table A.5: OLS and IV Estimates — Mortality (Heart Attacks)

	OLS	OLS	First Stage	Reduced Form	IV	IV
	(1)	(2)	(3)	(4)	(5)	(6)
Smoke	0.012	0.009	−0.040***	0.038	−0.002	0.143*
	(0.071)	(0.070)	(0.013)	(0.040)	(0.031)	(0.080)
Medicaid	−0.031	−0.034			0.094	0.116
	(0.072)	(0.074)			(0.116)	(0.120)
Smoke × Medicaid	−0.007	−0.006				−0.047*
	(0.023)	(0.023)				(0.025)
Z			0.347***	0.038		
			(0.093)	(0.038)		
Smoke × Z				−0.025*		
				(0.015)		
Included controls:						
Year FE	✓	✓	✓	✓	✓	✓
County-month-of-year FE	✓	✓	✓	✓	✓	✓
Weather controls		✓	✓	✓	✓	✓
Socioeconomic controls		✓	✓	✓	✓	✓
Observations	4,176	4,176	4,176	4,176	4,176	4,176
Adjusted R^2	0.917	0.917	0.978	0.917	0.917	0.916

Note: *p<0.1; **p<0.05; ***p<0.01. The dependent variable in columns (1)-(3) and (5)-(7) is the log of the count of heart attack deaths in county j and month t of year y. The dependent variable in column (4) is the Medicaid coverage rate (the total number of Medicaid enrollees divided by total population) in county j and month t of year y. Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across county-months (equal to 15.4% coverage) or greater. Standard errors are clustered by county.

County-Level Correlation

Figure A.2: First Stage Correlation (County-Level)

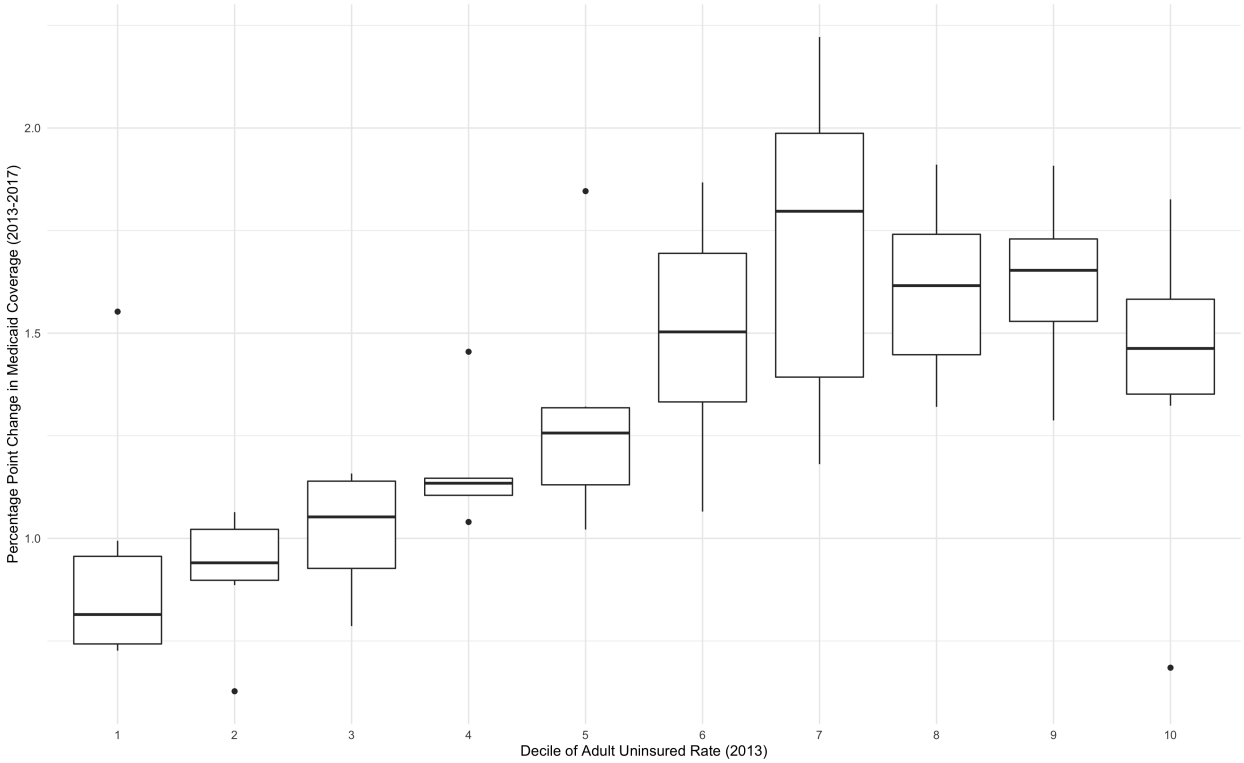

Notes: Figure shows the unconditional distributions of the percentage point change in County-level Medicaid coverage from 2013 to 2017 for each decile of the exogenous instrument (the adult uninsured rate in 2013). There are 58 total counties in California.

A.6 Double-IV Model

In this section, I present results in which $Smoke_{jty}$ is used as an exogenous instrument for mean $PM_{2.5}$ in the ZIP-week or county-month, rather than used as a reduced form measure of air pollution exposure. Consistent with the main text, $Smoke_{jty}$ is constructed as a binary variable representing one standard deviation or greater of smoke coverage across the ZIP-week or county-month, equal to 17.2% and 15.4%, respectively.

Table A.6: Double-IV Estimates — ER Visits (Selected Conditions)

	Selected Conditions				Selected Placebos		
	Asthma	All Resp.	Hypertension	Heart Attack	All Conditions	Diabetes	Appendicitis
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mean $PM_{2.5}$	0.033	0.014	−0.027	0.005	0.009	−0.036	−0.002
	(0.025)	(0.017)	(0.022)	(0.016)	(0.011)	(0.023)	(0.016)
Medicaid	0.003	0.050	−0.143	0.009	0.052	−0.011	0.032
	(0.165)	(0.128)	(0.209)	(0.084)	(0.111)	(0.152)	(0.096)
Mean $PM_{2.5}$ × Medicaid	−0.009	−0.005	0.004	−0.001	−0.004	0.010	0.001
	(0.008)	(0.005)	(0.006)	(0.004)	(0.004)	(0.007)	(0.005)
Observations	44,926	44,926	44,926	44,926	44,926	44,926	44,926
Adjusted R^2	0.746	0.929	0.505	0.349	0.982	0.639	0.303

Notes: *p<0.1; **p<0.05; ***p<0.01. The dependent variable is the log of the count of ER admissions for the stated condition in ZIP j and week t of year y. $PM_{2.5}$ is measured in standard units of $\mu g/m^3$. Medicaid is instrumented for using the rate of uninsured adults in the ZIP code in 2013 (as in the ER visits models in the main text), and Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. $PM_{2.5}$ is instrumented for with a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across ZIP-weeks (equal to 17.2% coverage) or greater. All models include all of the fixed effects and controls listed in Table 1.4 of the main text. Standard errors are clustered by ZIP.

Table A.7: Double-IV Estimates — Mortality (Selected Conditions)

	Selected Conditions				Selected Placebos		
	Asthma	All Resp.	Hypertension	Heart Attack	All Conditions	Diabetes	Appendicitis
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mean $PM_{2.5}$	0.028	0.025	0.011	0.067	−0.010	0.024	0.014
	(0.047)	(0.046)	(0.056)	(0.063)	(0.025)	(0.053)	(0.019)
Medicaid	0.042	0.085	0.019	0.360	−0.058	−0.066	0.044
	(0.193)	(0.192)	(0.264)	(0.293)	(0.121)	(0.293)	(0.080)
Mean $PM_{2.5}$ × Medicaid	−0.004	−0.007	−0.002	−0.020	0.006	−0.003	−0.003
	(0.014)	(0.014)	(0.017)	(0.020)	(0.009)	(0.016)	(0.006)
Observations	3,710	3,710	3,710	3,710	3,710	3,710	3,710
Adjusted R^2	0.924	0.950	0.915	0.909	0.989	0.894	0.111

Notes: *p<0.1; **p<0.05; ***p<0.01. The dependent variable is the log of the count of deaths for the stated condition in county j and week t of year y. $PM_{2.5}$ is measured in standard units of $\mu g/m^3$. Medicaid is instrumented for using the rate of uninsured adults in the county in 2013 (as in the mortality models in the main text), and Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. $PM_{2.5}$ is instrumented for with a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across county-months (equal to 15.4% coverage) or greater. All models include county-month-of-year and year fixed effects and weather and socioeconomic controls. Standard errors are clustered by county.

A.6.1 IV Estimates: ER Visits (Monitor Sample)

This table repeats the analysis of Table 1.5 in the main text, using only the ZIP codes that have official EPA monitors for $PM_{2.5}$. This equals 165 ZIP codes out of approximately 1,750 total in California.

Table A.8: IV Estimates (Monitor Sample) — Mortality (Selected Conditions)

	Selected Conditions				Selected Placebos		
	Asthma	All Resp.	Hypertension	Heart Attack	All Conditions	Diabetes	Appendicitis
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Smoke	0.064	0.038	−0.026	0.004	0.024	−0.065*	−0.007
	(0.042)	(0.030)	(0.037)	(0.028)	(0.022)	(0.037)	(0.031)
Medicaid	−0.122	−0.020	−0.081	−0.002	0.006	0.119	0.045
	(0.214)	(0.131)	(0.154)	(0.061)	(0.137)	(0.168)	(0.071)
Smoke × Medicaid	−0.017	−0.014*	−0.002	0.001	−0.009	0.015	0.003
	(0.013)	(0.008)	(0.010)	(0.008)	(0.007)	(0.010)	(0.009)
Observations	44,926	44,926	44,926	44,926	44,926	44,926	44,926
Adjusted R^2	0.757	0.933	0.519	0.349	0.983	0.655	0.304

Notes: *p<0.1; **p<0.05; ***p<0.01. Sample includes only ZIPs that have at least one regulatory grade monitor for $PM_{2.5}$. The dependent variable is the log of the count of ER visits for the stated condition in ZIP j and week t of year y. Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across ZIP-weeks (equal to 17.2% coverage) or greater. All models include ZIP-week-of-year and year fixed effects and weather and socioeconomic controls. Standard errors are clustered by ZIP.

A.6.2 IV Estimates: Mortality (Monitor Sample)

This table repeats the analysis of Table 1.6 in the main text, using only the counties that have official EPA monitors for $PM_{2.5}$. This equals 52 counties out of 58 total in California.

Table A.9: IV Estimates (Monitor Sample) — Mortality (Selected Conditions)

	Selected Conditions				Selected Placebos		
	Asthma	All Resp.	Hypertension	Heart Attack	All Conditions	Diabetes	Appendicitis
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Smoke	0.050	0.050	0.020	0.134	−0.022	0.043	0.027
	(0.077)	(0.089)	(0.107)	(0.090)	(0.047)	(0.086)	(0.031)
Medicaid	−0.007	−0.003	−0.005	0.104	0.012	−0.106	0.003
	(0.058)	(0.052)	(0.100)	(0.105)	(0.051)	(0.127)	(0.022)
Smoke × Medicaid	−0.004	−0.013	−0.003	−0.038	0.012	−0.003	−0.005
	(0.021)	(0.026)	(0.032)	(0.027)	(0.016)	(0.026)	(0.009)
Observations	3,710	3,710	3,710	3,710	3,710	3,710	3,710
Adjusted R^2	0.927	0.951	0.915	0.918	0.990	0.896	0.139

Notes: *p<0.1; **p<0.05; ***p<0.01. Sample includes only counties that have at least one regulatory grade monitor for $PM_{2.5}$.The dependent variable is the log of the count of deaths for the stated condition in county j and month t of year y. Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across county-months (equal to 15.4% coverage) or greater. All models include county-month-of-year and year fixed effects and weather and socioeconomic controls. Standard errors are clustered by county.

A.7 Impact Estimates

Table A.10: Estimated Benefits — Avoided Smoke-Related ER Visits and Deaths

	(a) ER Visits (Asthma)				(b) Mortality (Heart Attacks)			
		Δ Medicaid Rate				Δ Medicaid Rate		
Year	Actual	+ 5pp	+ 10pp	+ 20pp	Actual	+ 5pp	+ 10pp	+ 20pp
2012	−3,756	−68	−136	−272	219	−4	−8	−15
2013	−20,229	−230	−460	−920	1,247	−8	−16	−31
2014	−7,491	−102	−203	−406	507	−5	−11	−21
2015	−15,476	−184	−368	−735	932	−8	−17	−33
2016	−35,381	−255	−510	−1,019	2,044	−10	−20	−40
2017	−35,938	−257	−514	−1,029	2,018	−12	−24	−49
Total	118,271	−1,096	−2,191	−4,381	6,967	−47	−96	−189

Note: Counterfactual estimates of avoided ER visits and deaths are statewide totals, and are based on predictions using the estimates of τ_{IV} in Table 1.5 (ER visits) and Table 1.6 (mortality) for asthma and heart attacks, respectively. The estimates reflect actual and predicted health outcomes *only* during periods defined as smoky, i.e. having at least one standard deviation or greater of overall smoke coverage during the ZIP-week (for ER visits) or county-month (for mortality). There were relatively fewer wildfires (and thus relatively fewer associated smoke plumes) in 2012 and 2014 as compared to the other years in the sample.

A.8 Robustness Checks

A.8.1 ER Visits (County-Month)

This table repeats the analysis of ER visits but instead uses a county-by-month aggregation rather than a ZIP-by-week aggregation, to facilitate comparisons with the mortality analysis which can only be conducted at the county-month level due to data limitations.

Table A.11: IV Estimates — ER Visits (Selected Conditions)

	Selected Conditions				Selected Placebos		
	Asthma	All Resp.	Hypertension	Heart Attack	All Conditions	Diabetes	Appendicitis
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Smoke	0.054	0.078**	−0.138*	−0.008	0.010	−0.112	0.059
	(0.068)	(0.036)	(0.081)	(0.055)	(0.019)	(0.070)	(0.092)
Medicaid	−0.018	−0.026	−0.155	0.035	0.036	−0.004	0.237**
	(0.169)	(0.097)	(0.160)	(0.094)	(0.036)	(0.079)	(0.116)
Smoke × Medicaid	−0.028	−0.026**	0.033	−0.004	−0.006	0.030	−0.012
	(0.022)	(0.011)	(0.025)	(0.018)	(0.005)	(0.023)	(0.031)
Observations	4,176	4,176	4,176	4,176	4,176	4,176	4,176
Adjusted R^2	0.977	0.991	0.969	0.961	0.998	0.979	0.954

Notes: *p<0.1; **p<0.05; ***p<0.01. The dependent variable is the log of the count of ER visits for the stated condition in county j and month t of year y. Medicaid-related estimates have been scaled to reflect a 10pp increase in Medicaid coverage. Smoke is a binary variable indicating smoke coverage of *at least* one standard deviation of the smoke distribution across county-months (equal to 15.4% coverage) or greater. All models include county-month-of-year and year fixed effects and weather and socioeconomic controls. Standard errors are clustered by county.

Appendix B

Supplemental Material for Chapter 2

B.1 Main Results: Excluding California

Table B.1: The effect of upwind emissions on monitors' $PM_{2.5}$ levels under CAIR Excluding California

Dependent variable: Monitor-average $PM_{2.5}$ (month of sample)		
	(1) Pooled	**(2) Annual Standard**
$\widehat{PM_{2.5}}$	1.255***	
	(0.205)	
$(\widehat{PM_{2.5}})$ × Far, Below		1.079***
		(0.210)
$(\widehat{PM_{2.5}})$ × Near, Below		1.334***
		(0.210)
$(\widehat{PM_{2.5}})$ × Near, Above		1.433***
		(0.204)
$(\widehat{PM_{2.5}})$ × Far, Above		1.261***
		(0.268)
$(\widehat{PM_{2.5}})$ × Far, Below × CAIR		0.294
		(0.264)
$(\widehat{PM_{2.5}})$ × Near, Below × CAIR		0.409
		(0.279)
$(\widehat{PM_{2.5}})$ × Near, Above × CAIR		0.227
		(0.262)
$(\widehat{PM_{2.5}})$ × Far, Above × CAIR		-0.555
		(0.354)
Fixed effects		
Monitor	Yes	Yes
Month by NERC region	Yes	Yes
Observations	57,681	57,681

*p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the monitor-month, limited to only one monitor per county (selected monitors represent the "marginal" monitor, defined as the monitor that determined the county's annual design value for $PM_{2.5}$ in the respective year). The dependent variable is the monthly average $PM_{2.5}$ concentration at the given monitor.$\widehat{PM_{2.5}}$ refers to InMAP-based estimates of the transport of $PM_{2.5}$ from upwind emissions at electricity-generating units to the monitor. Distance indicators denote the distance of the monitor's county from the Annual $PM_{2.5}$ NAAQS of 15.0 $\mu g/m^3$ in the base period of CAIR (2007-2008). Based on a bandwidth of 1.0 $\mu g/m^3$, there are 619 monitors categorized as *Far Below*, 95 as *Near Below*, 47 as *Near Above*, and 9 as *Far Above*. We two-way cluster errors at the same level as the specifications' fixed effects. The data sample for this set of regressions covers 2005-2015.

Table B.2: The effect of upwind emissions on monitors' $PM_{2.5}$ levels under CSAPR/MATS Excluding California

Dependent variable: Monitor-average $PM_{2.5}$ (month of sample)		
	(1) Pooled	**(2) Annual Standard**
$\widehat{PM_{2.5}}$	1.259***	
	(0.177)	
$(\widehat{PM_{2.5}})$ × Far, Below		1.231***
		(0.192)
$(\widehat{PM_{2.5}})$ × Near, Below		1.182***
		(0.178)
$(\widehat{PM_{2.5}})$ × Near, Above		1.192***
		(0.234)
$(\widehat{PM_{2.5}})$ × Far, Above		1.256***
		(0.173)
$(\widehat{PM_{2.5}})$ × Far, Below × CAIR		0.199
		(0.261)
$(\widehat{PM_{2.5}})$ × Near, Below × CAIR		0.178
		(0.270)
$(\widehat{PM_{2.5}})$ × Near, Above × CAIR		0.107
		(0.285)
$(\widehat{PM_{2.5}})$ × Far, Above × CAIR		-0.417*
		(0.246)
$(\widehat{PM_{2.5}})$ × Far, Below × CSAPR/MATS		-0.160
		(0.516)
$(\widehat{PM_{2.5}})$ × Near, Below × CSAPR/MATS		-0.498
		(0.507)
$(\widehat{PM_{2.5}})$ × Near, Above × CSAPR/MATS		-1.466**
		(0.683)
$(\widehat{PM_{2.5}})$ × Far, Above × CSAPR/MATS		-0.760*
		(0.453)
Fixed effects		
Monitor	Yes	Yes
Month by NERC region	Yes	Yes
Observations	79,274	79,274

*p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the monitor-month, limited to only one monitor per county (selected monitors represent the "marginal" monitor, defined as the monitor that determined the county's annual design value for $PM_{2.5}$ in the respective year). The outcome variable is the monthly average $PM_{2.5}$ concentration at the given monitor. $\widehat{PM_{2.5}}$ refers to InMAP-based estimates of the transport of $PM_{2.5}$ from upwind emissions at electricity-generating units to the monitor. Distance indicators denote the distance of the monitor's county from the Annual $PM_{2.5}$ NAAQS of 12.0 $\mu g/m^3$ in the base period of CSAPR/MATS (2013-2014). Based on a bandwidth of 1.0 $\mu g/m^3$, there are 753 monitors categorized as *Far Below*, 69 as *Near Below*, 19 as *Near Above*, and 2 as *Far Above*. We two-way cluster errors at the same level as the specifications' fixed effects. The data sample for this set of regressions covers 2005-2019.

Table B.3: The effect of upwind emissions on monitors' $PM_{2.5}$ levels under CSAPR/MATS Excluding California, Above Category Collapsed

Dependent variable: Monitor-average $PM_{2.5}$ (month of sample)		
	(1) Pooled	**(2) Annual Standard**
$\widehat{PM_{2.5}}$	1.259***	
	(0.177)	
$(\widehat{PM_{2.5}})$ × Far, Below		1.234***
		(0.193)
$(\widehat{PM_{2.5}})$ × Near, Below		1.184***
		(0.179)
$(\widehat{PM_{2.5}})$ × Above		1.219***
		(0.205)
$(\widehat{PM_{2.5}})$ × Far, Below × CAIR		0.193
		(0.261)
$(\widehat{PM_{2.5}})$ × Near, Below × CAIR		0.173
		(0.270)
$(\widehat{PM_{2.5}})$ × Above × CAIR		-0.034
		(0.285)
$(\widehat{PM_{2.5}})$ × Far, Below × CSAPR/MATS		-0.156
		(0.516)
$(\widehat{PM_{2.5}})$ × Near, Below × CSAPR/MATS		-0.494
		(0.507)
$(\widehat{PM_{2.5}})$ × Above × CSAPR/MATS		-1.302**
		(0.616)
Fixed effects		
Monitor	Yes	Yes
Month by NERC region	Yes	Yes
Observations	79,274	79,274

*p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the monitor-month, limited to only one monitor per county (selected monitors represent the "marginal" monitor, defined as the monitor that determined the county's annual design value for $PM_{2.5}$ in the respective year). The outcome variable is the monthly average $PM_{2.5}$ concentration at the given monitor. $\widehat{PM_{2.5}}$ refers to InMAP-based estimates of the transport of $PM_{2.5}$ from upwind emissions at electricity-generating units to the monitor. Distance indicators denote the distance of the monitor's county from the Annual $PM_{2.5}$ NAAQS of 12.0 $\mu g/m^3$ in the base period of CSAPR/MATS (2013-2014). Based on a bandwidth of 1.0 $\mu g/m^3$, there are 753 monitors categorized as *Far Below*, 69 as *Near Below*, and 21 as *Above*. We two-way cluster errors at the same level as the specifications' fixed effects. The data sample for this set of regressions covers 2005-2019.

B.2 Placebo Tests: Mortality

Table B.4: The effect of changes in upwind emissions on changes in county mortality under CAIR

	(1) Diabetes	(2) Liver	(3) Congen.
$(\widehat{\text{PM}_{2.5}})$ × Far, Below	1.739	0.343	0.073
	(1.847)	(0.956)	(0.333)
$(\widehat{\text{PM}_{2.5}})$ × Near, Below	0.416	−0.040	0.341
	(1.731)	(0.896)	(0.312)
$(\widehat{\text{PM}_{2.5}})$ × Near, Above	0.678	0.009	−0.025
	(1.838)	(0.952)	(0.332)
$(\widehat{\text{PM}_{2.5}})$ × Far, Above	1.251	−0.826	0.439
	(3.075)	(1.592)	(0.555)
Fixed effects			
NERC Region	Yes	Yes	Yes
Observations	438	438	438

*p<0.1; **p<0.05; ***p<0.01. The outcome variables is the change in the county-level mortality rate per 100,000 residents for the stated health condition from the pre-period to the post-period of the policy. For CAIR, this change is calculated as the difference between the 2013-2014 average rate and the 2007-2008 average rate. $\Delta\widehat{PM_{2.5}}$ refers to the change in average InMAP-based exposure from 2007-2008 to 2013-2014. Distance indicators denote the distance of the county from the Annual $\text{PM}_{2.5}$ NAAQS of 15.0 $\mu g/m^3$ in the base period of CAIR (2007-2008). We cluster errors by NERC region. *Liver* includes chronic liver disease and cirrhosis. *Congen.* refers to congenital diseases.

Table B.5: The effect of changes in upwind emissions on changes in county mortality under CSAPR/MATS

	(1) Diabetes	(2) Liver	(3) Congen.
$(\widehat{\text{PM}_{2.5}})$ × Far, Below	−8.071*	−1.007	−1.031
	(4.633)	(2.059)	(0.864)
$(\widehat{\text{PM}_{2.5}})$ × Near, Below	−6.760	−0.189	−1.788*
	(5.324)	(2.366)	(0.993)
$(\widehat{\text{PM}_{2.5}})$ × Near, Above	−2.344	1.103	−2.727*
	(8.147)	(3.621)	(1.520)
$(\widehat{\text{PM}_{2.5}})$ × Far, Above	−0.696	−1.724	−1.729
	(15.126)	(6.723)	(2.822)
Fixed effects			
NERC Region	Yes	Yes	Yes
Observations	438	438	438

*p<0.1; **p<0.05; ***p<0.01. The outcome variables is the change in the county-level mortality rate per 100,000 residents for the stated health condition from the pre-period to the post-period of the policy. For CSAPR/MATS, this change is calculated as the difference between the 2017-2018 average rate and the 2013-2014 average rate. $\Delta\widehat{PM_{2.5}}$ refers to the change in average InMAP-based exposure from 2013-2014 to 2017-2018. Distance indicators denote the distance of the county from the Annual $\text{PM}_{2.5}$ NAAQS of 15.0 $\mu g/m^3$ in the base period of CSAPR/MATS (2013-2014). We cluster errors by NERC region. *Liver* includes chronic liver disease and cirrhosis. *Congen.* refers to congenital diseases.

Appendix C

Supplemental Materials for Chapter 3

C.1 Placebo Test

Table C.1: OLS Estimates — ER Visits (Bone Fractures)

	(1) 0-17	(2) 0-17
$< 40°$	0.063	0.085
	(1.01)	(1.20)
$40° - 49°$	0.056	0.036
	(0.58)	(0.44)
$50° - 59°$	0.027	0.026
	(0.52)	(0.41)
$70° - 79°$	0.095*	0.070
	(1.97)	(1.29)
$80° - 89°$	0.076	0.054
	(1.03)	(0.79)
$> 90°$	-0.009	-0.006
	(-0.07)	(-0.04)
$< 40° \times$ Medicaid		0.033
		(0.52)
$40° - 49° \times$ Medicaid		-0.056
		(-0.77)
$50° - 59° \times$ Medicaid		-0.042
		(-0.47)
$60° - 69° \times$ Medicaid		0.008
		(0.13)
$70° - 79° \times$ Medicaid		-0.042
		(-1.22)
$80° - 89° \times$ Medicaid		-0.019
		(-0.90)
$> 90° \times$ Medicaid		-0.080
		(-0.77)
Observations	10,160	20,320

Note: t-statistics in parentheses. *p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the county-quarter in column (1) and the county-quarter-Medicaid status in column (2). The outcome variable is the log of the count of inpatient admissions that originated in the Emergency Department and contain a bone fractures diagnosis. Total precipitation, per capita income, and the logged poverty rate are also controlled for in each regression. Standard errors are clustered by county.

C.2 Robustness

Table C.2: OLS Estimates — ER Visits (Severe Conditions)

	(1) 0-64	(2) 0-17	(3) 18-44	(4) 45-64
$< 40°$	-0.0008	-0.0002	0.0004	-0.0011
	(-1.16)	(-0.12)	(0.57)	(-1.42)
$40° - 49°$	-0.0002	-0.0021	0.0020**	-0.0007
	(-0.41)	(-1.16)	(2.39)	(-1.16)
$50° - 59°$	0.0001	-0.0021	0.0010	-0.0001
	(0.31)	(-1.63)	(1.52)	(-0.28)
$70° - 79°$	0.0001	-0.0003	0.0006	-0.0001
	(0.26)	(-0.17)	(1.40)	(-0.15)
$80° - 89°$	0.0002	-0.0002	0.0001	0.0001
	(0.33)	(-0.15)	(0.18)	(0.23)
$> 90°$	0.0008	-0.0027	0.0008	0.0007
	(0.55)	(-1.53)	(0.51)	(0.46)
$< 40° \times$ Medicaid	-0.0000	-0.0002	-0.0001	-0.0000
	(-0.00)	(-0.16)	(-0.03)	(-0.00)
$40° - 49° \times$ Medicaid	-0.0029***	-0.0026**	-0.0008	-0.0035***
	(-2.99)	(-2.30)	(-0.67)	(-3.00)
$50° - 59° \times$ Medicaid	-0.0010	0.0020**	-0.0014*	-0.0010
	(-1.03)	(2.57)	(-1.94)	(-0.86)
$60° - 69° \times$ Medicaid	-0.0021***	-0.0003	0.0000	-0.0027***
	(-4.32)	(-0.34)	(0.03)	(-4.84)
$70° - 79° \times$ Medicaid	-0.0014***	-0.0000	-0.0010***	-0.0015***
	(-5.61)	(-0.03)	(-2.76)	(-5.54)
$80° - 89° \times$ Medicaid	-0.0016***	-0.0000	-0.0007**	-0.0019***
	(-4.09)	(-0.03)	(-2.32)	(-4.10)
$> 90° \times$ Medicaid	-0.0025	0.0040***	-0.0026	-0.0026
	(-1.40)	(3.14)	(-1.08)	(-1.39)
Observations	20,320	20,320	20,320	20,320

Note: t-statistics in parentheses. *p<0.1; **p<0.05; ***p<0.01. The unit of analysis is the county-quarter in column (1) and the county-quarter-Medicaid status in column (2). The outcome variable is the log of the count of inpatient admissions that originated in the Emergency Department and contain a severe respiratory and/or cardiovascular-related diagnosis (cardiac dysarythmias and congestive heart failure). Total precipitation, per capita income, and the logged poverty rate are also controlled for in each regression. Standard errors are clustered by county.