

UC Davis

UC Davis Electronic Theses and Dissertations

Title

Essays on Environmental Economics: Air Pollution and Cognitive Decline

Permalink

<https://escholarship.org/uc/item/3w009649>

ISBN

9798189226199

Author

Firsova, Irina

Publication Date

2026-06-25

Peer reviewed|Thesis/dissertation

Essays on Environmental Economics: Air Pollution and Cognitive Decline

By

IRINA FIRSOVA
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Economics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Erich Muehlegger, Chair

Marianne Bitler

Jenna Stearns

Committee in Charge

2026

To my Grandpa, Seryozha.
I miss you, and I love you, and I wish you were here to see.

Acknowledgments

This dissertation, in one way or another, is a product of so many people. There are no words in the English language to describe how grateful I am to everyone who showed me kindness and support along the way.

First of all, I want to thank Erich and Marianne for being absolutely the best advisors one could ask for. This journey is already difficult as it is, even in the most stable of conditions. On top of that, multiple curveballs have been thrown at me throughout. There were so many prolonged periods of time when I just couldn't take it anymore and wanted to give up. I genuinely think I would have, had it not been for your support. I am so so grateful for your kindness and for your patience, and for never letting me doubt I can make it through. I appreciate you teaching me how to be a better researcher. Even more so, I appreciate you teaching me the importance of guiding with compassion. I hope I can be that person for someone else in the future. Thank you for everything.

I also want to thank Janine Wilson for teaching me everything I know about teaching. It would have taken me much longer to figure out that this is what I want to do, had it not been for your trust and encouragement. Your approach to being an educator is so incredible, and I'd be lucky to become half as good at it as you are. Thanks for believing in me, and I'll do my best.

I spent a lot of nights questioning everything, and will likely spend many more doing the same thing. This must be the nature of leaving everything behind. It was really easy to visualize what I left. It was really hard to visualize what I left for. In the grand scheme of things, none of it really mattered. For obvious reasons, what I left behind wasn't there anymore, and I never had anywhere to go back to. Questioning where my home is was another common activity of mine. In retrospect,

a very silly thing to question, because the answer turned out to be very simple. How lucky am I to have two places to call home.

I don't know, life's crazy! If you told me seven years ago that this is how it would all turn out, I would not have believed any of it. And yet, here we are. It is what it is, and what it's not, it's not. And that's just how things are. For the most part, I think I made it through just fine, but I do have a lot of people to thank for it.

Growing up, for many reasons, some in my control, some out of it, I found it hard to feel like I belong. I will forever be grateful to everyone in Davis for helping me break this cycle. Maybe the real PhD was the friends we made along the way...

To my friends, of which I have many. Soqui, Carlos, Reem, Sebastian, Stephen, Alle, Beau, Colter, Bruno. I used to watch 'Friends' growing up, and always thought this is not something that happens to real people, or to people like me. Thanks for proving me wrong, fellas. I could not have asked for a better Clique. Such an incredible group of friends IS What I've Been Looking For. Will cherish this forever.

I'll miss coming home and having Soqui and Carlos there. I'll miss going to the movies with Stephen and Sebastian. I'll miss having Vinho Verde with Reem. I'll miss having parties at ABC house. I'm putting it all down because one day I'll want to know I was there. I'll miss it all so much.

Thank you to Ellen, for all the good times and all the support. To Beryl, for being the kindest friend to me. To Joe and Pedro, for helping me not go crazy through the pandemic. To Gabriel and Rafael, for all the fun in San Francisco. To Joaquin, for being a proud Regal Unlimited member

with me. It does pay for itself if you go twice a month... To Yeji and Pierro, for all the cute dinners. I'm sure I'm missing someone. And it's okay!

To Masha, thank you for all the help, time and sense of direction on my way to here. To Alya, for being my connection to home and all the good times in New York. To Julia, for everything. I would not have been brave enough if it wasn't for you. I wish I could tell you that.

Бабушка Люда, Бабушка Галя, Дедушка Толя! Спасибо вам за любовь и поддержку, даже на таком большом расстоянии я никогда об этом не забываю. Дедушка Серёжа, надеюсь ты мной гордишься. Я тебя люблю и очень скучаю.

To my parents. None of it would be possible without your unconditional and constant support. I will forever appreciate every single thing, and all the effort you put into just being there, whether it's at my best or at my worst. I would not be where I am if it wasn't for you. Knowing you got my back makes every single difficult thing so much easier. Thank you for your love and acceptance. I miss you all the time. I love you, Mom!! I love you, Dad!!

To Laurin, the love of my life. Thank you for making it all make sense. The future is looking very bright. I'm gonna need the darkest of glasses! Having you by my side is the greatest privilege, and you make my life better in every dimension. All the scary things in life don't feel scary anymore. I missed you for twenty eight years, and I love you one billion.

To my sister, Kristina. To my dear, dear sister. Every day it's driving me crazy that we are so far apart. I'm praying for the day when we are finally not. Everything good in me and around me is better because I have you. Everything bad in me and around me is less bad for the same reason. You are and always will be the most important person in my life. The only person I'd run

through an airport for. I just would not be who I am, or have any sense of self without you. I am so so proud of you, always. I don't know what I did in my previous life to deserve having you as my sister. Thanks for being my forever best friend. I will never take it for granted.

I will end this on one of my favorite quotes by one of my favorite Davids.

'I really enjoyed forgetting. When I first come to a place, I notice all the little details. I notice the way the sky looks. The color of white paper. The way people walk. Doorknobs. Everything. Then I get used to the place and I don't notice those things anymore. So only by forgetting can I see the place again as it really is.'

I can't wait to forget.

Abstract

Governments have long recognized air pollution as a negative externality with significant economic and health costs. Fine particulate matter has been linked to various negative health outcomes (cardiovascular disease, respiratory illness, premature death). More recently, researchers have found that pollution also impairs cognition and mental well-being. These consequences are harder to observe and are usually absent from standard cost-benefit analyses. This dissertation argues that as a result, the true social costs of air pollution are likely underestimated. This dissertation examines different ways in which air pollution can lead to cognitive decline. I study the effect of pollution exposure on fatal traffic crashes, crash frequency and severity, and mental health. The main challenge to estimate causal effects of air quality is that pollution exposure is not randomly assigned. I address this using instrumental variable strategies, exploiting exogenous variation in wind direction and attainment status, to establish causal estimates of how pollution affects cognitive outcomes.

The first chapter estimates the effect of fine particulate matter ($PM_{2.5}$) on fatal car crashes across the United States. Using daily variation in wind direction as an instrument for local pollution, I find that a $1 \mu g/m^3$ increase in $PM_{2.5}$ raises daily fatal accidents by 1.34 percent. Extending the analysis to hourly data, the effect rises to 1.97 percent for pollution in the hour preceding a crash. These results suggest that short-term exposure impairs cognitive function in real time. These estimates imply that reducing $PM_{2.5}$ by $1 \mu g/m^3$ could prevent roughly 550 fatalities annually, equivalent to about \$4.06 billion in savings based on the value of a statistical life.

The second chapter extends this analysis using comprehensive crash data from Chicago. Unlike national records that capture only fatal crashes, the Chicago data include every reported incident,

allowing me to separate pollution's effect on crash frequency from its effect on severity. Instrumenting for $PM_{2.5}$ with wind direction and taking advantage of Chicago's geography, where Lake Michigan and concentrated industrial sources lie in different directions, I find that a $1 \mu g/m^3$ increase in $PM_{2.5}$ raises total crashes by 0.78 percent and fatal crashes by 9.05 percent. Individual-level analysis further shows that higher pollution increases the severity of injuries conditional on a crash occurring, indicating that pollution affects not only how often crashes happen but how severe they are.

The third chapter examines whether cleaner air improves mental health outcomes. I exploit regulatory variation created by the 1970 Clean Air Act Amendments, using total suspended particulate nonattainment status as an instrument for changes in pollution between 1970 and 1980, within a first-differenced design that absorbs time-invariant county characteristics. Two-stage least squares estimates of the effect on county-level suicide rates are small and statistically insignificant across specifications, and heterogeneity analysis by race and gender yields uniformly insignificant results. The population-weighted estimates are positive and significant, suggesting that effect may be concentrated in large urban counties.

These essays show that the social costs of air pollution extend beyond respiratory and cardiovascular health. The strong and consistent effects on traffic fatalities suggest a cognitive channel that should be included in cost-benefit analyses. The burden of pollution, and the benefits of reducing it, may be substantially larger than commonly recognized.

Contents

1 Does Air Pollution Affect the Likelihood of a Fatal Car Crash? Evidence from the US	1
1.1 Introduction	2
1.2 Background on PM (particulate matter)	9
1.3 Data	13
1.3.1 Data on car accidents	14
1.3.2 Data on air pollution	15
1.3.3 Data on atmospheric conditions	18
1.4 Identification and Empirical strategy	19
1.5 Results for daily analysis	23
1.6 Heterogeneity and Nonlinearity	28
1.7 Hourly estimates	30
1.8 Discussion	33

2	Air Pollution and Traffic Crashes: Evidence from Chicago	56
2.1	Introduction	57
2.2	Data	64
2.2.1	Crash Data	64
2.2.2	Air Pollution Data	67
2.2.3	Weather Data	69
2.3	Effects of Pollution on Crashes	70
2.3.1	First Stage Results	73
2.3.2	OLS Results	73
2.3.3	IV Results	74
2.4	Effects of Pollution on Crash Severity	76
2.5	Conclusion	83
3	Effect of the Clean Air Act (1970) on Mental Health Outcomes	102
3.1	Introduction	103
3.2	Institutional background	108

3.2.1	Background on the Clean Air Act	108
3.2.2	Background on PM (particulate matter)	110
3.3	Data	111
3.3.1	Air Pollution Data	111
3.3.2	Suicide Rate Data	114
3.3.3	Descriptive Statistics	115
3.3.4	Empirical Framework	116
3.3.5	Control Variables	118
3.4	Identification and Empirical Strategy	119
3.4.1	First-Difference Specification	119
3.4.2	OLS Results	120
3.4.3	Instrumental Variable Strategy	121
3.4.4	2SLS Results	123
3.5	Heterogeneity Analysis	125
3.5.1	Motivation	125

3.5.2	Racial Heterogeneity	125
3.5.3	Gender Heterogeneity	127
3.5.4	Results	127
3.6	Limitations	129
3.7	Conclusion	131
	References	143

Chapter 1:

Does Air Pollution Affect the Likelihood of a Fatal Car Crash? Evidence from the US

Irina Firsova

University of California, Davis

Abstract

This paper examines the impact of fine particulate matter ($\text{PM}_{2.5}$) on fatal car accidents across the United States, using instrumental variable (IV) analysis with wind direction as an instrument for air pollution. I find that a $1 \mu\text{g}/\text{m}^3$ increase in daily $\text{PM}_{2.5}$ levels is associated with a 1.34% increase in daily fatal accidents after accounting for weather conditions and fixed effects. Extending the analysis to hourly data, I find that a $1 \mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ levels in the hour prior to an accident results in a 1.97% increase in fatal crashes, suggesting that short-term pollution exposure impairs cognitive function and increases accident risk in real time. Reducing $\text{PM}_{2.5}$ by $1 \mu\text{g}/\text{m}^3$ could prevent approximately 550 fatalities annually, leading to an estimated economic savings of \$4.06 billion based on the value of a statistical life (VSL). These findings contribute to the literature on the cognitive effects of air pollution and suggest that the social costs of pollution may be underestimated.

1.1 Introduction

Motor vehicle crashes are one of the leading causes of death in the United States, killing more than 38,000 people annually. Car accidents are a preventable cause of death, and there is a large body of literature examining how various policies affect traffic fatalities. While in Europe the rate of traffic fatalities has been decreasing, in the United States it is on the rise. The European Commission published a report suggesting that during 2019-2022 the motor vehicle fatality rate decreased by 10% in the EU. Simultaneously, according to CDC National Center for Health Statistics, from year 2019 to 2021 there was a 16% increase in motor vehicle fatalities in the United States. Properly understanding what parameters affect the motor vehicle fatality rate can be crucial when designing a policy aimed at reducing it. Most literature focuses on policies to which drivers can directly respond and adjust their behavior on the road, such as mandatory seat-belt use ([Cohen and Einav \(2003\)](#)), drunk driving laws ([Eisenberg \(2003\)](#)), and speed limits ([Gupta, Johari, Jashami, and Savolainen \(2022\)](#)). Before a policy becomes law, there is a time period during which drivers may prepare to use motor vehicles under a new policy. This paper contributes to a growing literature examining drivers' response to exogenous shocks, including studies on traffic volume shocks such as COVID-19 safer-at-home policies ([Brodeur, Cook, and Wright \(2021\)](#)), environmental factors such as seasonal pollen ([Danagouliau and Deza \(2024\)](#)), and economic shocks such as stock market fluctuations ([Giulietti, Tonin, and Vlassopoulos \(2020\)](#)). It is important to know, from a social planner's perspective, if there is a common and unexpected shock that can affect drivers' ability to operate a motor vehicle. This knowledge can be useful for designing effective interventions, such as real-time alerts or temporary restrictions on

high-risk driving conditions. Quantifying the risks associated with such shocks allows one to conduct a more accurate cost-benefit analysis. Measures like pollution advisories or dynamic speed limits could help mitigate accident risks and save lives, reducing the economic burden of traffic fatalities.

In this paper I estimate the magnitude of the short-run (recent) effect of exposure to $PM_{2.5}$, examining exposure windows of one day and one hour, on fatal car accidents in the United States. Using air quality data from the Environmental Protection Agency (EPA) and traffic fatality data from the Fatality Analysis Reporting System (FARS), I aggregate pollution and accident data to the county-day level over the period from 1999 to 2013. The main concern when estimating the causal impact of air pollution on traffic fatalities is that exposure to air pollution is not randomly assigned, which can lead to biased estimates. Selection bias may arise because pollution monitors are disproportionately located in densely populated urban areas, where both traffic and accident rates are higher, potentially overstating the true relationship. Reverse causation is less likely to be a concern here, as it is implausible that traffic accidents themselves cause meaningful changes in ambient pollution levels. Omitted variable bias is the most serious issue. Factors such as local economic activity or seasonal traffic patterns may simultaneously drive both pollution levels and accident rates, leading to upward bias in OLS estimates. Taken together, these concerns suggest that OLS estimates of the pollution-accident relationship cannot be given a causal interpretation. To address this, I use an instrumental variable (IV) approach, exploiting daily variations in wind direction as an exogenous source of variation in pollution levels. After accounting for weather conditions and fixed effects, I assert the exclusion restriction: changes in a county's wind direction affect fatal car crashes only through their impact on air pollution, and not

through any other channel. By using variation in wind patterns, I isolate the effects of air pollution on traffic mortality rates. This exclusion restriction is plausible because wind direction is determined by large-scale atmospheric dynamics and is unlikely to directly influence driving behavior or accident risk independently of the pollution it carries. The relevance condition is satisfied. Wind direction is a robust predictor of local $PM_{2.5}$ levels, with a first-stage F-statistic of 370.5, well above conventional thresholds for instrument strength. A key benefit of the instrumental variable methodology is that it eliminates the need to isolate the source of pollution. This is particularly important in the context of air quality studies, where multiple emission sources—such as industrial facilities, vehicle traffic, and natural events like wildfires—can simultaneously contribute to pollution levels. By using exogenous variation, such as changes in wind direction, the IV approach allows me to isolate the impact of pollution on traffic accidents, without needing to take the source of pollution into account. In addition, the inclusion of county, state-by-month, and month-by-year fixed effects accounts for systematic variation in traffic volume across days and regions, ensuring that the estimated effect of pollution is not confounded by days when both traffic and pollution happen to be high.

My main air pollutant of interest is $PM_{2.5}$. Particulate matter (PM) consists of tiny solid and liquid particles suspended in the air, and it is regulated by the EPA. Fine particles under 2.5 micrometers in diameter, including combustion particles, metals, and organic compounds, are called $PM_{2.5}$. Not only do these smaller particles enter the lungs but they also reach the bloodstream. Natural sources of $PM_{2.5}$ include volcanic eruptions and wildfires, while anthropogenic sources come from fossil fuel combustion in power plants, industries, and vehicles. $PM_{2.5}$ can remain airborne for extended periods and travel long distances

and penetrate buildings unless the buildings are equipped with air filtering system. Long-term exposure to $PM_{2.5}$ is associated with premature death, particularly in individuals with chronic heart or lung diseases, as well as impaired lung development in children.

I find that a one microgram per cubic meter ($\mu\text{g}/\text{m}^3$) increase in $PM_{2.5}$ levels leads to a 0.0009 increase in daily fatalities per county. This finding suggests that a one ($\mu\text{g}/\text{m}^3$ increase in $PM_{2.5}$) corresponds to approximately a 1.34 percent increase in fatal car accidents, or roughly 550 additional deaths annually across the United States. Using the value of a statistical life (VSL, estimated at \$7.4 million in 2006 by the EPA; [EPA \(2006\)](#)), which is commonly used by social planners to quantify the economic value of reducing mortality risks, this translates to \$4.06 billion. To put this figure in context, the EPA’s own estimates place the total annual benefits of Clean Air Act programs at approximately \$110 billion ([U.S. Environmental Protection Agency, 1999](#)), suggesting that the road safety channel accounts for nearly 4% of pollution’s total estimated economic costs. Moreover, the GDP losses from the 2008–2009 recession totaled approximately \$1.2 trillion over two years ([Congressional Budget Office, 2012](#)). The annual road safety costs of air pollution identified here represent 0.3% of that, which is a recurring burden that accumulates without policy intervention. These results hold after controlling for temporal variation in atmospheric conditions such as wind speed, temperature, and precipitation, as well as geographic fixed effects to account for seasonal and regional variations in both air quality and traffic patterns. This paper provides a stronger rationale for adopting stricter air quality standards and pollution reduction policies, since true costs of pollution might be underestimated.

In addition to the daily analysis, I extend the study to examine how immediate exposure

to pollution affects road safety using hourly data from 2010 to 2013. The hourly results reveal that a 1 $\mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ levels leads to a 1.97% increase in fatal car crashes. These findings highlight the cognitive effects of pollution on drivers' decision-making and reaction times. Including day-of-week fixed effects in the hourly analysis is important as it allows me to capture weekday and weekend driving patterns that influence accident likelihood. This more granular approach demonstrates how pollution's impact on real-time driving safety is not uniform throughout the day but fluctuates with variations in traffic intensity and driving behavior.

The relationship between the daily and hourly results reveals how pollution affects drivers over different time frames. The daily analysis captures average pollution exposure across all hours, including hours when most people are not driving, which may dilute the measured effect. The hourly analysis measures pollution in the hour immediately preceding an accident, capturing exposure when drivers are on the road. This helps explain why the hourly estimate is larger. At the same time, daily county-level pollution remains a meaningful measure for residents, since approximately 72 percent of US workers live and work in the same county ([McKenzie, 2013](#)). It implies that daily ambient pollution levels reflect the environment in which most drivers both live and commute. Together, the daily and hourly findings suggest that both accumulated and immediate exposure to pollution matter. However, immediate exposure may be especially detrimental in high-risk situations.

In this paper I find that short-term (recent) exposure to particulate matter leads to a statistically significant increase in road fatalities. These results suggest that the benefits of environmental policy and ameliorating air pollution are underestimated, as improving air

quality can result in fewer accidental deaths on the road, which in turn would lead to lower public health expenditures. From a social planner’s perspective, there is a tradeoff between health costs and economic benefits when implementing a tax on air pollution. Understanding the health costs of air pollution is essential for environmental policy proposals. If health costs are not estimated correctly, it will negatively impact the efficiency of environmental policy.

The effects of air pollution on health (Kim, Jahan, and Kabir (2013), Fiordelisi, Piscitelli, Trimarco, Coscioni, Iaccarino, and Sorriento (2017), Segalowitz (2008)), productivity (Carson, Koundouri, and Nauges (2011), Hanna and Oliva (2015), He, Liu, and Salvo (2019)), and crime (Herrnstadt, Heyes, Muehlegger, and Saberian (2021)) are well documented. Air pollution is a significant public policy issue with major negative consequences. Multiple studies have shown that exposure to air pollution leads to adverse health outcomes, mainly through respiratory and cardiovascular diseases (Kim et al. (2013), Fiordelisi et al. (2017)). Poor air quality imposes a heavy burden on public health expenditures (Segalowitz (2008)) as chronic diseases and disabilities caused by pollution lead to a large increase in healthcare costs.

The cognitive effects of $PM_{2.5}$, including impaired decision-making and reduced attention, are discussed in detail in Section 2.

This paper builds on the methodology from Deryugina et al. (2019), using their instrumental variable (IV) strategy to establish causality between air pollution and adverse health or safety outcomes. The validity of an IV approach requires that daily variations in wind direction are strongly correlated with $PM_{2.5}$ levels, and that wind direction affects fatal car crashes only through its effect on pollution. Deryugina et al. have established that changes in

wind direction are a good predictor of local pollution concentrations. The exclusion restriction requires that wind direction affects fatal car crashes only through its effect on pollution levels, and not through any other channel. After controlling for weather conditions and fixed effects, I argue this assumption is plausible: wind direction is unlikely to directly influence driving behavior or accident risk independently of the pollution it carries. I am able to isolate the short-term effects of particulate matter on traffic safety, reducing potential bias in my estimates that might arise from confounding factors like driving patterns or regional economic conditions.

This paper contributes to growing literature on negative effects of air pollution on cognitive functioning. The role of air quality in traffic safety has been less explored. By estimating the extent to which short-term exposure to $PM_{2.5}$ can increase fatal car crashes, my findings highlight the broader societal costs of pollution and its influence on cognitive function in high-stakes, real-time decision-making scenarios. It has been established that policies regarding mandatory seat-belt use ([Cohen and Einav \(2003\)](#)), drunk driving ([Eisenberg \(2003\)](#)), and speed limits ([Gupta et al. \(2022\)](#)) significantly reduced road fatalities. A closely related paper is [Sager \(2019\)](#). It examines the impact of air pollution on road safety in the United Kingdom between 2009 and 2014. Using temperature inversions as a source of exogenous variation, they find that a $1\mu g/m^3$ increase in $PM_{2.5}$ levels leads to a 0.3–0.6% increase in the number of vehicles involved in road accidents per day. My results are consistent with those findings. Car travel is substantially more prevalent in the United States than in the United Kingdom, making the road safety consequences of pollution particularly important in the US context. Most closely related is [Burton and Roach \(2026\)](#), which examines how exposure to particulate matter pollution impairs cognition and increases fatal car crashes

in the United States using daily data. This paper extends that analysis to hourly observations, which reveals patterns that daily data cannot capture. In particular, the hourly analysis shows that the effect of pollution on fatal crashes is strongest during evening rush hour, when traffic is dense and driving demands are highest, and that immediate exposure in the hour prior to an accident has a larger estimated effect than daily average exposure. This suggests that cognitive impairment from pollution is most consequential precisely when drivers face the greatest navigational demands.

The remainder of this paper is organized as follows. Section 2 provides background information on particulate matter (PM), describing the different types of PM and their sources, as well as the health risks associated with exposure. Section 3 outlines the data sources used in this paper. In Section 4, I present the identification strategy and empirical methodology. Section 5 discusses the daily results. Section 6 examines heterogeneity and nonlinearity in the relationship between $PM_{2.5}$ and fatal crashes. In Section 7 I present my results for hourly level analysis. Finally, Section 8 concludes the paper by summarizing the findings and their implications for public policy.

1.2 Background on PM (particulate matter)

Particulate matter is a term used to describe microscopic particles of solid and liquid matter suspended in the air. There are two types of particles that are regulated by the Environmental Protection Agency (EPA) in the US. Particulate matter with diameter of between 2.5 and 10 micrometers are called “coarse particulates,” while PM less than 2.5 in diameter

are called “fine particulates.”

PM_{10} particles are inhalable particles less than 10 micrometers in diameter. For example, dust, pollen, and mold are PM_{10} particles. These particles are small enough to get into the lungs. Exposures to PM_{10} have been linked to the worsening of respiratory diseases, such as asthma, leading to hospitalization and emergency department visits.

$PM_{2.5}$ particles are inhalable particles less than 2.5 micrometers in diameter. For example, combustion particles, metal, and organic compounds can be $PM_{2.5}$ particles. These particles are small enough to get into the lungs and enter the bloodstream. The main natural sources of $PM_{2.5}$ are volcanic eruptions and wildfires. As for anthropogenic sources, they include fossil fuel combustion from power plants, industries, and automobiles. They can stay in the air for a long time and travel for hundreds of miles. They can enter buildings: hence, many people are easily exposed to it. Long-term exposure to $PM_{2.5}$ has been associated with premature death, particularly in people who have chronic heart or lung diseases, and reduced lung function growth in children.

$PM_{2.5}$ is considered to be the most dangerous form of air pollution from a public health perspective, with epidemiological studies providing evidence of a causal relationship between long-term $PM_{2.5}$ exposure and premature death. Estimates suggest that $PM_{2.5}$ is responsible for approximately 52,000 premature deaths in the United States annually ([Lelieveld, Evans, Fnais, Giannadaki, and Pozzer, 2015](#)). Once $PM_{2.5}$ enters the lungs and bloodstream, it has been linked to a range of cardiovascular and respiratory diseases, including heart attacks, strokes, chronic obstructive pulmonary disease, and lung cancer ([Pope, Burnett, Thun, Calle, Krewski, Ito, and Thurston \(2002\)](#); [Brook, Rajagopalan, Pope III, Brook, Bhatnagar,](#)

Diez-Roux, Holguin, Hong, Luepker, Mittleman et al. (2010); Dockery, Pope, Xu, Spengler, Ware, Fay, Ferris, and Speizer (1993)). Long-term exposure to high levels of $PM_{2.5}$ is also associated with premature death, particularly in vulnerable populations, such as those with existing heart or lung conditions Hoek, Krishnan, Beelen, Peters, Ostro, Brunekreef, and Kaufman (2013); Pope and Dockery (2006). $PM_{2.5}$ can remain suspended in the air for a long time and travel significant distances - even thousands of miles - from the source. If pollution originates in industrial regions, it can still affect rural areas or areas much further away from it. Additionally, $PM_{2.5}$ particles can penetrate indoor environments, making it difficult to avoid exposure, even when indoors. The negative effects of being exposed to $PM_{2.5}$ go beyond cardiovascular issues. Recent studies show that $PM_{2.5}$ can also affect the cognitive functioning of the human brain (Weuve, Puett, Schwartz, Yanosky, Laden, and Grodstein (2012)) . Particulate matter can cross the blood-brain barrier, which can cause inflammation in the brain, leading to problems with memory, decision-making and ability to concentrate (Calderon-Garciduenas, Franco-Lira, Torres-Jardon, Henriquez-Roldan, Barragan-Mejia, Valencia-Salazar, Gonzales-Maciel, Reynoso-Robles, Villarreal-Calderon, and Reed (2008)). This is especially concerning for activities like driving, where full attention and quick thinking are crucial.

Beyond cardiovascular and respiratory effects, $PM_{2.5}$ exposure has been shown to impair cognitive functioning through multiple channels. The most intuitive is reduced work attendance due to illness, but a less obvious and perhaps more consequential channel is a direct decline in cognitive performance among exposed individuals (Hausman, Ostro, and Wise (1984)). Overall, the negative effects of air pollution on the cognitive system are well-studied.

Given that exposure to air pollution has been shown to impair cognitive functioning, it is reasonable to expect a negative impact on tasks requiring focus and quick decision-making, such as driving. Driving a motor vehicle demands sustained attention and quick reflexes, all of which may deteriorate under the influence of particulate matter. If air pollution negatively impacts cognitive function, drivers are more prone to errors and delayed reactions on the road. This could lead to an increase in traffic accidents.

Studies have also linked air pollution to impaired real-time decision-making, showing increased errors among MLB umpires ([Archsmith, Heyes, and Saberian \(2018\)](#)). While these errors often occur in relatively low-risk settings, driving involves rapid decisions where even small mistakes can lead to fatal accidents. This makes the decline in cognitive function caused by pollution exposure particularly dangerous on the road. Research on the effects of air pollution on mortality has predominantly focused on internal causes of death and hospitalizations ([Deryugina, Heutel, Miller, Molitor, and Reif \(2019\)](#)). However, the consequences of air pollution extend beyond these severe health outcomes, posing immediate risks to public safety.

$PM_{2.5}$ is the primary pollution measure in this paper for two reasons. First, as discussed above, $PM_{2.5}$ is the most harmful form of particulate matter. Its small size allows it to penetrate deep into the lungs, enter the bloodstream, cross the blood-brain barrier, and infiltrate indoor environments, making exposure difficult to avoid ([Pope et al. \(2002\)](#); [Brook et al. \(2010\)](#)). Coarser particles such as PM_{10} are filtered more effectively by the respiratory system and do not carry the same systemic health risks. Second, $PM_{2.5}$ is the most consistently monitored pollutant. The EPA has provided comprehensive data since 1999, whereas data

for other pollutants such as nitrogen dioxide (NO_2), sulfur dioxide (SO_2), ozone (O_3), carbon monoxide (CO), and lead are much less consistently available over time and across locations. Third, $PM_{2.5}$ is invisible and odorless, meaning that drivers are unlikely to observe or respond behaviorally to changes in its concentration. This is important for identification. If drivers could detect pollution and adjust their behavior accordingly, for example, by driving more cautiously or avoiding the road altogether, this would confound the estimated effect of pollution on accidents. Coarser or more visible pollutants, such as smoke or smog, do not share this property.

1.3 Data

This paper combines three datasets to estimate the causal effect of $PM_{2.5}$ on fatal car accidents: traffic fatality data from the Fatality Analysis Reporting System (FARS), air quality data from the Environmental Protection Agency (EPA), and atmospheric data from the North American Regional Reanalysis (NARR). I aggregate all three sources to the county-day level over the period 1999 to 2013 for the daily analysis, and to the county-hour level over 2010 to 2013 for the hourly analysis. The dependent variable is the number of fatal car crashes per county per day (or hour). $PM_{2.5}$ is the primary independent variable of interest. Wind direction, wind speed, temperature, and precipitation serve as atmospheric controls and as the basis for the instrumental variable strategy. Summary statistics for the daily analysis are presented in Table 1.1. The average number of fatal crashes per county per day is 0.067, reflecting the relative rarity of fatal accidents at the county level. The average $PM_{2.5}$ concentration is $10.60 \mu\text{g}/\text{m}^3$, which is above the current EPA annual standard of $9 \mu\text{g}/\text{m}^3$.

Summary statistics for the hourly analysis are presented in Table 1.2. The average number of fatal crashes per county per hour is 0.00358, consistent with the daily mean scaled to an hourly level. The average $PM_{2.5}$ concentration is slightly lower at $9.60 \mu\text{g}/\text{m}^3$, reflecting the more restricted 2010–2013 sample period.

1.3.1 Data on car accidents

Data on fatal car accidents is obtained from the Fatality Analysis Reporting System (FARS). It records every fatal car accident on public roads in the United States. I aggregate FARS data to represent a county-daily measure. On average there is one fatal car accident in a county in a day. Figure 1a represents the raw trend in car accidents over my time period of interest. On average there is from 19,000 to 24,000 fatal car accidents in a year in the United States in the areas where air pollution is monitored.

The trend in accidents shows a steady rise from January, starting at around 2,000 accidents, and peaking in July with approximately 3,000 accidents. After July, the total number of accidents gradually declines, with the year ending in December at just over 2,500 accidents. The increase in accidents from winter to mid-summer likely reflects several external factors. For example, summer months see a surge in travel, with more people on the road for vacations, contributing to the rise in accidents from about 2,000 in January to 3,000 in July. Additionally, summer hazards like heavy rains or extreme heat in certain regions may also elevate the risk of accidents. The decline after July could be attributed to a reduction in travel as the vacation season ends, causing the number of accidents to drop from 3,000 in July to around 2,500 by December. This decrease may also reflect improved weather conditions

in some regions as summer storms subside. Furthermore, as people return to their regular routines and schools reopen in late August and September, traffic patterns shift. Fewer leisure drivers and more predictable commuting behavior may contribute to the post-July decline in accidents. These seasonal patterns in raw accident counts motivate the inclusion of state-by-month and month-by-year fixed effects in the empirical analysis. Combined with county fixed effects, these controls absorb the seasonal and regional variation visible in the raw data, so that identification comes from day-to-day variation in wind-driven pollution within a county and season, rather than from cross-seasonal or cross-regional comparisons (Figure 1.6).

1.3.2 Data on air pollution

Air pollution data for pollutants regulated by the Clean Air Act ($PM_{2.5}$, ozone, carbon monoxide, sulfur dioxide, nitrogen dioxide) and PM_{10} are taken from the Environmental Protection Agency (EPA) Air Quality database. I am using dataset from [Deryugina et al. \(2019\)](#) that contains information on all pollutants at the monitor level. All available monitor readings within counties are averaged to obtain county-level measures. My pollutant of interest is $PM_{2.5}$ and I use data at the daily level. For the daily analysis, I analyze data from 1999 to 2013, as the comprehensive data for $PM_{2.5}$ is available from 1999. For the hourly analysis covering 2010 to 2013, I use hourly $PM_{2.5}$ readings from the same EPA monitors. The hourly analysis covers a shorter period than the daily analysis because consistent hourly $PM_{2.5}$ monitor readings are only available from 2010 onward. Prior to 2010, EPA monitors reported primarily daily averages, with hourly readings being sparse and unevenly distributed

across monitors and counties. Each monitor reports a concentration reading for every hour of the day, and these readings are averaged across all monitors within a county to produce county-hour measures of $PM_{2.5}$. This finer temporal resolution allows me to match pollution exposure to the specific hour preceding each accident, rather than relying on a daily average that includes hours when most people are not driving. Figures 1b and 2b show the average monthly and yearly $PM_{2.5}$ concentrations across the United States over the 1999–2013 period. The average $PM_{2.5}$ levels steadily decline over time, from $13.6 \mu g/m^3$ (micro-grams per cubic meter) in year 1999 to $8.1 \mu g/m^3$ in year 2013. The number of pollution monitors remained approximately the same since year 2001. Beyond the long-run decline, $PM_{2.5}$ levels also vary at shorter time horizons. For example, across seasons, across days within a season, and even across hours within a day, driven by various factors. It is important to note that according to [Sullivan, Krupnick et al. \(2018\)](#), counties can strategically place pollution monitors in cleaner areas, which introduces measurement error. Monitor readings may systematically understate true exposure levels faced by residents. The instrumental variable approach addresses this measurement error directly, since wind direction is determined atmospherically and is orthogonal to the location decisions of pollution monitors. Changes in wind direction shift pollution levels across the entire county regardless of where monitors are placed. It means that the IV isolates variation in true exposure rather than variation in measured exposure (Figure 1.7).

Although only about a third of U.S. counties are covered by pollution monitors, I am able to capture 64% of all car accidents during the study period. Monitors are mainly located in more populated and urban areas where both traffic and pollution levels are higher. I

miss some rural areas due to sparse monitoring of pollution.¹ To the extent that the effect of pollution on fatal crashes also exists in unmonitored rural areas, my estimates likely represent a lower bound on the true effect.²

On average, $PM_{2.5}$ levels start relatively high in January, with concentrations around $11 \mu\text{g}/\text{m}^3$, then steadily decline to a low of approximately $8 \mu\text{g}/\text{m}^3$ by April. However, concentrations rise sharply through the spring, reaching a peak of nearly $12 \mu\text{g}/\text{m}^3$ in July. After this summer high, the concentration drops again, hitting its lowest point of the year—just over $9 \mu\text{g}/\text{m}^3$.

This pattern can be explained by a combination of seasonal factors. The mid-year rise in $PM_{2.5}$ concentrations, especially in June and July, likely corresponds to higher traffic due to summer travel, and possibly wildfires in certain regions. These months tend to see more outdoor activities, which often contribute to higher pollution levels. On the other hand, the elevated $PM_{2.5}$ levels in January and December could be driven by increased heating during the colder months, with more homes and businesses burning fuel for warmth, leading to higher emissions. The sharp decline in spring and late fall reflects a period of reduced heating and potentially less travel or industrial activity, helping to lower pollution levels during those months.

¹I rely on ground-level EPA monitor data rather than satellite-based or other modeled exposure measures. Satellite data offer broader geographic coverage but introduce their own measurement error, particularly in areas with frequent cloud cover or complex terrain, and are not available at the hourly resolution required for the hourly analysis.

²Rural drivers may be less exposed to high $PM_{2.5}$ concentrations on average, but may also be more vulnerable due to longer emergency response times and less access to healthcare. Omitting rural areas could bias the estimated effect downward.

1.3.3 Data on atmospheric conditions

Data on wind direction and wind speed for 1999–2013 are available from the North American Regional Reanalysis (NARR) daily reanalysis data. Wind direction is defined as the direction the wind is blowing from. For example, if wind direction is SW, it means the wind is blowing from the South-West. Wind speed and direction are reported as u- and v-vector components on a 32×32 kilometer grid, interpolated to each pollution monitor location, converted to wind speed and direction, and then averaged to the county-day level. Clean data on wind direction are obtained from [Deryugina et al. \(2019\)](#). For the hourly analysis, all weather controls are assigned at the county-hour level using hourly readings from the EPA.

Temperature and precipitation data come from [Schlenker and Roberts \(2009\)](#), who produce a daily weather grid using PRISM and weather station data. Total daily precipitation and daily maximum and minimum temperatures are reported for each point on a 2.5×2.5 mile grid covering the contiguous US for 1999–2013. These grid-point measures are then averaged across all grid points within a county to produce county-day measures of tmax, tmin, and total precipitation. Rather than entering tmax and tmin as continuous variables, I follow [Deryugina et al. \(2019\)](#) and bin them into indicators. There are 17 bins for daily maximum temperature and 17 bins for daily minimum temperature, ranging from -15°C (5°F) or below up to 30°C (86°F) or above, with 15 intermediate bins each spanning 3°C (5.4°F). This binned indicator approach is more flexible than a polynomial specification, as it imposes no assumptions about the functional form of the temperature-accident relationship and allows for nonlinear effects at any point in the temperature distribution.

1.4 Identification and Empirical strategy

I am estimating short run effect of exposure to fine particulate matter on the likelihood of getting in a fatal car crash. The relationship can be described by the following equation

$$Y_{dmyc} = \beta PM_{2.5dmyc} + \theta X'_{dmyc} + \sigma_c + \sigma_d + \sigma_{my} + \sigma_{sm} + \epsilon_{dmyc} \quad (1.1)$$

where the dependent variable (Y_{dmyc}) is the number of fatal car accidents in county c , on day d , in month m in year y . β is the main coefficient of interest and captures the effect of a one $\mu\text{g}/\text{m}^3$ increase in daily $PM_{2.5}$ exposure on the number of fatal car accidents. A vector of time-varying control variables is represented by X_{dmyc} . I control for weather conditions because they affect both pollution levels and driving safety independently of each other. For example, rain reduces visibility and road traction, extreme temperatures alter driving behavior, and wind speed influences accident risk. Specifically, I include indicators for extremely high temperature levels (above 29°C) and extremely low temperatures (below 0°C), precipitation rates in deciles, and wind speed in miles per hour. To account for geographic differences in car traffic and air pollution, I include county (σ_c) fixed effects. State-by-month fixed effects ((σ_{sm})) control for any seasonal correlation between car accidents, wind direction and air pollution, as well as allowing for this correlation to vary by state. Month-by-year fixed effects (σ_{my}) account for common time-varying shocks, such as those induced by any environmental or car-related policy changes during the period of this study. Day-of-week fixed effects (σ_d) account for systematic differences in driving behavior across days of the week, with Monday through Friday (M–F) characterized by commuting traffic and

Saturday–Sunday by leisure travel. Standard errors are clustered at the county level. The interval I am considering for my analysis at daily level spans 1999 to 2013. My methodology is based on [Deryugina et al. \(2019\)](#).

OLS estimates are very likely to be biased for various reasons. First of all, exposure to fine particulate matter is not randomly assigned. Pollution monitors’ locations are not random, with higher monitor counts typically showing up in areas with higher population density. This non-random placement implies that monitor readings might not fully record actual exposure faced by all individuals in a county, leading to underestimated or overestimated levels of pollution exposure, depending on the local geography. Furthermore, people can drive through areas with different levels of pollution than where they live or are typically exposed. This creates a potential discrepancy between the monitor readings and the actual exposure of drivers, introducing measurement error and potentially leading to downwards bias in OLS estimates since measurement error is misrepresenting the actual exposure to pollution. Overall, the presence of measurement error and omitted variable bias is very likely. I exploit variation in air pollution due to changes in daily wind direction in order to estimate the causal effect of exposure to fine particulate matter on getting into a car accident. There are two key identifying assumptions in my analysis. The first, the exclusion restriction, requires that, controlling for weather conditions and fixed effects, wind direction affects fatal car crashes only through its effect on $PM_{2.5}$ levels, and not through any other channel. I argue this assumption is plausible. Wind direction is determined atmospherically and is unlikely to be correlated with local economic conditions or driving behavior. Unlike temperature or precipitation, wind direction does not systematically vary with industrial activity or traffic patterns, making it a credibly exogenous source of variation. One remaining

concern is that wind conditions could directly affect driving safety, for example, strong winds making vehicles harder to control. This concern is addressed by controlling for wind speed in the regression. Wind direction, conditional on wind speed, should affect accidents only through the pollution it transports. The second identifying assumption is relevance. The wind direction instruments must significantly move $PM_{2.5}$ levels, given the other controls. Unlike the exclusion restriction, this assumption is directly testable. It is confirmed by the first stage regression, which shows that wind direction is a strong predictor of local pollution concentrations, with a first-stage F-statistic of 370.5, well above conventional thresholds for instrument strength.

My first stage equation is defined the following way

$$PM2.5_{cdmy} = \sum_{g \in G} \sum_{b=0}^2 \beta_b^g 1[G_c = g] * WINDDIR_{cdmy}^{90b} + X'_{cdmy} + \gamma_c + \gamma_{ms} + \gamma_{my} + \epsilon_{cdmy} \quad (1.2)$$

where $PM2.5_{cdmy}$ represents the fine particulate matter levels in county c , on day d , in month m in year y . Variables $1[G_c = g] * WINDDIR_{cdmy}^{90b}$ are the excluded instruments. Wind directions are split into four 90-degree bins $[90b, 90b+90]$ such that $b \in \{0, 2\}$. Excluded reference bin is $[270, 360]$, which represents the West-Northwest wind direction. Results are robust to increasing the number of bins and coefficients are very similar if I change the range of b to $b \in \{0, 7\}$ splitting wind direction into 8 bins. For the purpose of reducing the computational burden of this model I am sticking to $b \in \{0, 2\}$. Vector of controls X'_{dmyc} and fixed effects are defined in equation (1).

Following [Deryugina et al. \(2019\)](#), I group pollution monitors into spatial clusters using

a k-means algorithm. While this step is not required for identification given the IV strategy, it serves to reduce measurement error arising from sparse monitor coverage and locally produced pollution. Specifically, monitor locations within counties are quite widespread, as a result, pollution-monitor readings in a county might misrepresent the actual average fine particulate matter exposure for county residents. In [Deryugina et al. \(2019\)](#), k-means cluster algorithm is used to classify all the pollution monitors in the United States into a hundred spatial groups based on their location. Grouping counties into spatial clusters based on proximity to monitors can help avoid the issue of sparse monitoring. Neighboring monitors are more likely to be assigned to the same group than the monitors far away from each other. On average each geographic group contains 21 monitors and 9 counties. Indicator for a county c being assigned to a group g from the set of monitor groups $\mathbf{G}$ is represented by the variable $1[G_c = g]$. This group division eliminates the issue of sparse monitor locations and limits the effect of wind direction in a county on its air pollution level to be the same within all counties assigned to a specific geographic area. Potential measurement error is also addressed through this approach. It is expected that the impact of the local sources of pollution emission will differ within a monitor group, based on the relative location of pollution monitor and pollution source. By splitting counties into different groups geographically, the impact of variation due to locally produced pollution is reduced. The most relevant example of locally produced pollution in this context is the pipe emissions from cars that can contain fine particulate matter. It is important to note here that another thing generating measurement error is that locally produced pollution is unlikely to reach all the people within the monitor group. However, the non-local pollution sources located on either side of the whole monitor group are expected to have the same effect on all monitors

within the group which makes non-local pollution emission likely to determine the variation in air pollution levels analyzed in equation (2), which is also helpful to reduce measurement error.

The coefficient of interest is β_b^g which represents the interaction between fine particulate matter and wind direction. It varies geographically between 100 assigned groups and between 4 different wind direction bins.

Wind direction as an instrument for air pollution is a commonly used instrument, and since it is a strong predictor of pollution levels, weak instruments are not a concern within this setting. My final specification has $N=3 \times 100=300$ instruments and my first stage F-statistics is sufficiently large and is equal to 370.5 (Figure 1.8, Figure 1.9).

1.5 Results for daily analysis

I begin by presenting OLS regression results, which are likely to be biased for the reasons outlined in Section 4. Results with and without controls are presented in Table 1.3. Results of the instrumental variable specification with and without controls are presented in Table 1.5. To facilitate direct comparison with the hourly analysis, IV estimates using the restricted 2010–2013 sample are presented in Table 1.6. The OLS estimates of effect of air pollution are statistically significant and show that a one $\mu g/m^3$ increase in the levels of fine particulate matter leads to an increase of 0.0002 in fatalities per capita. This effect is very small and represents a 0.02% increase. It becomes even smaller after controlling for atmospheric variables, signifying only a 0.001 increase in fatalities. Pollution exposure is often measured

with error, as monitors may not capture the true exposure experienced by all drivers across a county. This measurement error is unlikely to be purely classical white noise. It reflects two systematic challenges. First, sparse monitor coverage means that many county-days have no direct pollution reading, requiring interpolation or averaging across distant monitors. Second, monitors are disproportionately located in areas with worse air quality, meaning that measured pollution may systematically differ from true exposure depending on where within a county a driver is located. To the extent that this non-classical measurement error attenuates the OLS estimate, it biases the coefficient toward zero.

To alleviate potential downwards bias from measurement error and omitted variables, I am using IV specification. The IV results are statistically significant at a 1% level and suggest that a one $\mu\text{g}/\text{m}^3$ increase in the levels of fine particulate matter lead to increase of 0.0036 fatalities in a county per day. These results are discussed in detail below, with three models: (1) no controls or fixed effects, (2) with weather controls and county, month-year and state-month fixed effects, and (3) adding day-of-week fixed effects. The percentage change indicates that in the baseline model (column 1), a 1 $\mu\text{g}/\text{m}^3$ increase in $PM_{2.5}$ corresponds to a 5.43% increase in daily crashes. Adding controls for extremely high and low temperatures, wind speed and precipitation rate, as well as fixed effects (column 2) reduces the effect to 1.34%. Including Day-of-Week fixed effects (column 3) only slightly alters the effect to 1.35%. Given that driving patterns vary systematically across days of the week, with Monday through Friday (M–F) characterized by commuting traffic and Saturday–Sunday by leisure travel, day-of-week fixed effects are included as the preferred specification throughout the analysis. The stability of the coefficient across columns (2) and (3) provides reassurance that the results are not driven by day-of-week variation in either traffic or pollution. This

suggests that controlling for driving behavior throughout the week does not substantially affect the relationship between air pollution and crashes. The addition of weather controls, such as precipitation and temperature, significantly reduces the estimated effect. This suggests that part of the observed relationship in the baseline model could be attributed to weather factors. Fixed effects for counties, state-month, and month-year also improve model precision, reflected by the increase in the F-statistic from 370.5 to over 3,000 across the models. The stability of the $PM_{2.5}$ coefficient with and without the Day-of-Week fixed effects indicates that pollution’s impact on crashes is not driven by differences in weekday traffic patterns at the daily scale. Negative sign of coefficients for wind speed and precipitation rate is consistent with the literature. Higher wind speeds and precipitation may reduce crash frequency as drivers become more cautious, reduce speed, or avoid travel altogether during poor weather conditions. ([Knapp, Kroeger, Giese et al. \(2000\)](#)).

My final specification has $N=3 \times 100=300$ instruments and my first stage F-statistic is sufficiently large and is equal to 370.5, which is above the threshold to ensure no weak instrument bias. In other words, a one standard deviation increase in the levels of fine particulate matter leads to a 9% increase in daily fatal car crashes, which translates to approximately 2700 fatal crashes a year.

Fatal car crashes impose a substantial economic burden through medical costs. Reducing levels of $PM_{2.5}$ by one standard deviation could lead to a decrease in fatal car crash level by 2700 accidents a year on average. Using the value of a statistical life (VSL)—which is currently estimated at \$7.4 million by the U.S. Environmental Protection Agency ([EPA \(2006\)](#))—the potential economic cost of these additional 2700 fatalities amounts to \$19 billion

per year. This number represents the economic value of the lives lost due to traffic accidents driven by higher pollution levels and does not account for the further costs associated with non-fatal injuries, disabilities, or the long-term impacts on healthcare systems. The indirect effects, such as the strain on emergency response services, long-term healthcare costs, and psychological effects on victims' families, would add to the overall societal costs. These results highlight the benefits of environmental policy, not only through addressing health concerns but also through reducing motor vehicle fatalities. A well-designed policy could significantly improve both public health and road safety. Lowering $PM_{2.5}$ levels has the potential to save thousands of lives annually and avoid billions in economic costs.

Between 1999 and 2013, average concentrations of $PM_{2.5}$ fell by approximately 40%, from around $13.6 \mu\text{g}/\text{m}^3$ to $8.1 \mu\text{g}/\text{m}^3$ (Figure 2b). This decline could translate to significant reductions in fatal car accidents. With my findings showing that a $1 \mu\text{g}/\text{m}^3$ increase in $PM_{2.5}$ leads to a 1.34% increase in daily accidents, a $5.5 \mu\text{g}/\text{m}^3$ reduction (approximately 40%) would correspond to a 7.4% decrease in fatal crashes. Visual representation of this relationship.

Fatal accidents and $PM_{2.5}$ concentrations on average both peak in July. The rise in summer travelling could partially explain why we see more accidents and higher air pollution, as more vehicles on the road naturally lead to more emissions and a higher likelihood of traffic incidents. I argue that the cognitive effects of air pollution could also explain variation in fatal car accidents rate. When $PM_{2.5}$ concentrations are high—especially in busy, densely populated areas—it could impair drivers' attention, slow their reaction times, and impact decision-making. To ensure that I am not simply capturing the effect of summer's high

temperatures, I control for extreme heat (above 29°C), which could affect both pollution and accidents. Additionally, the use of month-by-year and state-by-month fixed effects accounts for broader seasonal patterns and region-specific differences to help isolate the effect.

Table 1.6 presents IV estimates using the restricted 2010–2013 sample, which covers the same period as the hourly analysis. Comparing the preferred specification (column 3) across the two samples, the estimated effect of a 1 $\mu\text{g}/\text{m}^3$ increase in $PM_{2.5}$ is 1.35% for the full 1999–2013 sample and 1.10% for the 2010–2013 subsample. The somewhat smaller estimate in the later period is consistent with the overall decline in $PM_{2.5}$ concentrations over time. As average pollution levels fell, the marginal effect of an additional unit of $PM_{2.5}$ may have diminished. The similarity of the two estimates provides reassurance that the main daily results are not driven by the earlier, higher-pollution years of the sample, and that the 2010–2013 period is representative of the full sample period.

Comparing the daily and hourly estimates over the same 2010–2013 period provides further insight into how the timing of pollution exposure matters. The preferred daily specification yields an effect of 1.10%, while the preferred hourly specification yields an effect of 1.97%. The larger hourly estimate suggests that pollution exposure in the specific hour preceding an accident is more predictive of crash risk than average daily exposure. This is consistent with the hypothesis that $PM_{2.5}$ impairs cognitive function in real time: a driver exposed to elevated pollution in the hour before getting on the road faces a more acute cognitive burden than one exposed to slightly higher average pollution over the course of a day. The gap between the two estimates also reflects the dilution effect of daily averaging, which includes hours when drivers are not on the road and pollution’s cognitive effects are

less consequential.

As a robustness check, Table 1.7 presents IV estimates using 8 wind direction bins rather than 4, over the same 2010–2013 sample. The preferred specification (column 3) has an effect of 1.26%, compared to 1.10% with 4 bins. The estimates are similar in magnitude and both statistically significant. It suggests that the results are not sensitive to the choice of wind direction binning. The first-stage F-statistics is large across all specifications. The slight increase in the estimated effect with 8 bins may reflect the finer resolution capturing more precise variation in pollution driven by wind direction.

1.6 Heterogeneity and Nonlinearity

To explore potential nonlinearity in the relationship between $PM_{2.5}$ and fatal car crashes, I split the sample at the median $PM_{2.5}$ concentration of $8.92 \mu\text{g}/\text{m}^3$. This threshold closely corresponds to the current EPA annual standard of $9 \mu\text{g}/\text{m}^3$. Tables 1.8 and 1.9 present IV estimates for the low and high pollution subsamples respectively. The effect of $PM_{2.5}$ on fatal crashes is statistically significant only in the high pollution subsample, where a $1 \mu\text{g}/\text{m}^3$ increase corresponds to a 1.23% increase in daily fatal crashes (column 3). In the low pollution subsample, the coefficient is positive but statistically indistinguishable from zero. This pattern suggests that $PM_{2.5}$ impairs driver cognition above a certain concentration level. The fact that this threshold aligns closely with the EPA annual standard has policy implications. Counties that already meet the standard may see limited road safety benefits from further reductions. But counties that remain above the standard would gain the most.

These findings suggest that pollution reduction policies targeting the highest-exposure areas and days would be most effective at reducing traffic fatalities.

Tables 1.10 and 1.11 present IV estimates by tercile and quartile of $PM_{2.5}$ concentration, providing a more granular picture of nonlinearity. The effect of $PM_{2.5}$ on fatal crashes is concentrated in the middle range of the pollution distribution and is statistically insignificant at both very low and very high pollution levels. In the tercile split, the effect is significant only in the second tercile, corresponding to a 8.77% increase per $\mu\text{g}/\text{m}^3$. In the quartile split, the effect is significant only in the third quartile, corresponding to a 17.61% increase. On extremely polluted days drivers may have reasons not to go outside. At very low pollution levels, $PM_{2.5}$ may be insufficient to meaningfully impair driver cognition. The range where the effect is strongest is roughly 7 to 14 $\mu\text{g}/\text{m}^3$.

As a further test of nonlinearity, Table 1.12 presents IV estimates of a polynomial specification in which both $PM_{2.5}$ and $PM_{2.5}^2$ are included as endogenous variables, instrumented using cardinal wind direction bins and their interactions. Across all three specifications, the coefficient on $PM_{2.5}$ is positive and the coefficient on $PM_{2.5}^2$ is negative. It is consistent with a concave relationship in which the marginal effect of pollution on crashes decreases at higher levels. In the preferred specification (column 3), the implied turning point is approximately 23.3 $\mu\text{g}/\text{m}^3$, which is above the EPA annual standard and above the range experienced by most US counties in the sample. This means that within the policy-relevant range of pollution levels, the marginal effect of $PM_{2.5}$ on fatal crashes is increasing.

1.7 Hourly estimates

This paper estimates the short-term effect of exposure to fine particulate matter (PM_{2.5}) on fatal car crashes in the United States, using hourly data from 2010 to 2013. To address potential endogeneity and measurement error, I employ an instrumental variable (IV) approach, where hourly variations in wind direction serve as an exogenous instrument for PM_{2.5}.

The equation I am estimating can be described as follows:

$$Y_{hmyc} = \beta \text{PM}_{hdmyc} + \theta X'_{hdmyc} + \sigma_c + \sigma_{sm} + \sigma_{my} + \sigma_d + \epsilon_{hdmyc}$$

where σ_c , σ_{sm} , σ_{my} , and σ_d denote county, state-by-month, month-by-year, and day-of-week fixed effects respectively, with day-of-week fixed effects (Monday through Friday versus Saturday–Sunday) included as the preferred specification.

The results of the IV regressions using hourly data are summarized in Table 1.13. Column (1) presents the baseline IV regression without any controls or fixed effects. The coefficient on PM_{2.5} is positive and statistically significant, indicating that an increase in air pollution correlates with more fatal car crashes. In Column (2), I add weather-related controls (precipitation, wind speed, extreme temperature) along with County, State-Month, and Month-Year fixed effects. This specification ensures that the results are not confounded by seasonal patterns, regional driving behaviors, or temperature extremes. The coefficient on

PM_{2.5} decreases slightly but remains statistically significant, suggesting an effect of air pollution on road safety after accounting for weather. Column (3) further adds Day-of-Week fixed effects to account for behavioral differences in driving patterns across weekdays and weekends. With these additional fixed effects, the PM_{2.5} coefficient remains significant.

The coefficient on PM_{2.5} in Column (3) suggests that a 1 $\mu\text{g}/\text{m}^3$ increase in hourly PM_{2.5} concentration leads to an increase of 0.000071 fatal crashes per hour. Given the average number of fatal crashes per hour is 0.0036, this effect translates into a 1.97% increase in fatal crashes. Similarly, in the baseline model (Column 1), the PM_{2.5} coefficient implies a 7.36% increase in fatal crashes per hour. This decrease in the estimated percentage effect across models suggests that some of the variation initially attributed to PM_{2.5} may be explained by weather conditions or driving patterns across different days of the week. The change in percentage effect from 0.44% to 1.97% reflects how hourly data captures patterns that can't be seen in daily data. Driving habits change throughout the day, for example, rush hours, late nights, and weekends all have distinct traffic patterns that influence accident risks. Including day-of-week fixed effects with hourly data can help to account for these fluctuations.

The negative coefficients on precipitation and wind speed align with previous findings in the literature. Higher wind speeds and precipitation may reduce crash frequency as drivers become more cautious, reduce speed, or avoid travel altogether during poor weather conditions. (Knapp et al. (2000)).

Table 1.14 presents IV estimates of the effect of PM_{2.5} on fatal car crashes separately for four time-of-day bins: morning rush hour (7–9am), midday (10am–3pm), evening rush hour

(4–7pm), and late night (8pm–6am). The effect of PM_{2.5} is largest and strongest during evening rush hour, where a 1 $\mu\text{g}/\text{m}^3$ increase in PM_{2.5} corresponds to a 3.76% increase in fatal crashes per hour ($p < 0.05$). The morning rush hour effect is positive and significant at the 10% level, corresponding to a 1.86% increase. By contrast, the midday effect is small (0.43%) and statistically indistinguishable from zero, and the late-night estimate is essentially zero (−0.04%). The ordering of effects is consistent with the pattern of evening rush hours having higher traffic density and more fatigued drivers, under these conditions cognitive impairment from pollution exposure would be especially strong.

This pattern is consistent with the hypothesis that air pollution impairs cognitive function when driving demands are greatest. During rush hours, drivers navigate dense traffic and operate under time pressure. Under these conditions even small declines in attention and reaction time can be fatal. Late-night driving involves sparse traffic and simpler road conditions, reducing the consequences of pollution-induced cognitive impairment even if the impairment itself is present.

Table 1.15 examines whether the effect of PM_{2.5} on fatal car crashes differs between weekdays and weekends. Both estimates are significant at the 10% level. Weekend driving shows a somewhat larger effect, with a 1 $\mu\text{g}/\text{m}^3$ increase in PM_{2.5} corresponding to a 1.48% increase in fatal crashes, compared to a 1.08% increase on weekdays. This pattern is consistent with a mechanism in which weekday commuters drive familiar routes on autopilot, potentially reducing their sensitivity to pollution-induced cognitive impairment, whereas weekend drivers are more likely to navigate unfamiliar roads or drive longer distances. However, these results should be interpreted cautiously. The key takeaway is that the effect of PM_{2.5} on

fatal crashes is present on both weekdays and weekends, suggesting the cognitive impairment mechanism is not specific to commuting behavior or workplace-related stress.

1.8 Discussion

In this paper, I look at how exposure to fine particulate matter $PM_{2.5}$ affects fatal car crashes in the United States. I estimate the effects of air pollution on traffic safety using hourly and daily data through an instrumental variable (IV) approach. Across six different model specifications, the findings show how pollution impacts accident risk in both the short and slightly longer term.

My main result is that a $1 \mu\text{g}/\text{m}^3$ increase in $PM_{2.5}$ leads to a 1.34% increase in daily accidents after accounting for weather conditions and fixed effects, down from an initial 5.43% with no controls. This decrease suggests that weather factors influence traffic safety independently of pollution.


I also find that being exposed to air pollution one hour prior to the accident increases the chances of getting into one. More specifically, a $1 \mu\text{g}/\text{m}^3$ increase in particulate matter levels results in a 1.97% increase in fatal accidents when controlling for day-of-week fixed effects. It is important to note that not taking into account day-of-week fixed effects reduces the estimate significantly, and the estimated effect is just 0.44%. It suggests that traffic patterns on weekdays and weekends mask some of the underlying impact of pollution that cannot be captured when data are aggregated at the daily level.

While accumulated exposure to poor air quality is important, the hourly analysis suggests that immediate exposure to pollution, like breathing in dirty air just before getting on the road, can impair drivers in real time, reducing their ability to make quick decisions. This explains why the effect of pollution on accidents becomes more pronounced when using granular data, especially with day-of-week fixed effects to capture behavioral differences.


These findings have significant policy implications. A reduction of $1 \mu\text{g}/\text{m}^3$ in $PM_{2.5}$ concentrations would prevent approximately 550 fatal car crashes annually across the United States. Using the EPA’s value of a statistical life of \$7.4 million, this translates to \$4.06 billion in annual economic savings through the road safety channel alone. To put this in perspective, the EPA estimates total annual benefits of Clean Air Act at approximately \$110 billion ([U.S. Environmental Protection Agency, 1999](#)). It means that the road safety channel identified here accounts for nearly 4% of the total estimated economic benefits of pollution reduction. Moreover, while the literature has documented substantial costs of air pollution through health outcomes such as cardiovascular disease and premature death ([Pope et al., 2002](#); [Brook et al., 2010](#)), reduced labor productivity ([Hanna and Oliva, 2015](#); [He et al., 2019](#)), and increased crime ([Herrnstadt et al., 2021](#)), the present findings suggest that these costs extend further to road safety. The \$4.06 billion estimate can be interpreted as a lower bound on the total social costs of $PM_{2.5}$ pollution, since it captures only one specific channel through which pollution harms society.

Although \$4.06 billion may appear modest relative to the total health-related social costs of air pollution, it is a very large number in absolute terms. For comparison, the GDP losses from the 2008–2009 recession totaled approximately \$1.2 trillion over two years ([Congres-](#)

sional Budget Office, 2012). The annual road safety costs of air pollution represent 0.3% of that figure, yet they recur every year without policy intervention. Policies such as real-time pollution advisories or stricter ambient air quality standards could help reduce this burden. The findings underscore that the social costs of air pollution are likely underestimated in existing cost-benefit analyses.




(a) Average Monthly Total Accidents in the US




(b) Average Monthly $PM_{2.5}$ Concentration in the US

Figure 1.1: These graphs show the monthly averages of total accidents and $PM_{2.5}$ concentrations across the United States for the years 1999–2013. The first graph illustrates the average total number of accidents per month, while the second graph shows the average concentration of $PM_{2.5}$ (fine particulate matter). The seasonal patterns in both graphs indicate potential correlations, particularly in the summer months, where higher pollution levels and increased accident rates coincide.



(a) Average Yearly Total Accidents in the US



(b) Average Yearly $PM_{2.5}$ Concentration in the US

Figure 1.2: These graphs display the yearly variations of total accidents and $PM_{2.5}$ concentrations across the United States. The first graph illustrates the total number of accidents per year, while the second graph shows the average concentration of $PM_{2.5}$.





Figure 1.3: **New York City: daily $PM_{2.5}$ concentrations and wind direction.** This figure shows first stage results for New York City area, with dependent variable being $PM_{2.5}$ levels, and independent variable being wind direction. Control variables and fixed effects are accounted for. Results are consistent with what is expected - if the wind blows from South-West (New Jersey), recorded air quality is low, and if the wind blows from the East (the ocean), recorded air quality is high.



Mean $PM_{2.5}$ in San Francisco, 1999–2015





Figure 1.4: **San Francisco: daily $PM_{2.5}$ concentrations and wind direction.** This figure shows first stage results for San Francisco area, with dependent variable being $PM_{2.5}$ levels, and independent variable being wind direction. Control variables and fixed effects are accounted for. Results are consistent with what is expected - if the wind blows from South-East (continental), recorded air quality is low, and if the wind blows from the West (the ocean), recorded air quality is high.

Average fatal accidents per day, 1999-2013



Average daily $PM_{2.5}$ levels, 1999-2013




Figure 1.5: Average fatal accidents per day and average daily $PM_{2.5}$ levels across US counties, 1999–2013. The top map displays the average number of fatal car accidents per day by county. The bottom map shows average daily $PM_{2.5}$ concentrations; counties in gray are not covered by pollution monitors.




Figure 1.6: Raw monthly averages of $PM_{2.5}$ concentrations and fatal crashes across the United States over the 1999–2013 period. Both peak in July, consistent with increased summer travel and higher pollution from seasonal factors such as wildfires. This motivates the inclusion of state-by-month fixed effects, which absorb this common seasonal variation.




Figure 1.7: Distribution of within-county $PM_{2.5}$ variation after removing county, state-month, month-year, and day-of-week fixed effects. The distribution confirms that day-to-day variation in pollution levels remains after accounting for all fixed effects. This variation is motivating the instrumental variable strategy.




Figure 1.8: This figure illustrates why controlling for seasonal patterns is important. After removing only county and year fixed effects, $PM_{2.5}$ residuals still display a seasonal pattern. This residual seasonal correlation motivates the inclusion of state-by-month fixed effects in the preferred specification.




Figure 1.9: Monthly averages of $PM_{2.5}$ and fatal crashes after removing the full set of fixed effects (county, state-month, month-year, and day-of-week). All systematic seasonal variation is absorbed, which leaves only day-to-day variation for identification.




Figure 1.10: Visual representation of the main result. It plots residualized fatal crashes against residualized $PM_{2.5}$ concentration, after partialling out county, state-month, month-year, and day-of-week fixed effects. The positive slope confirms the relationship identified in the IV analysis.




Figure 1.11: Estimated effect of $PM_{2.5}$ on fatal crashes from the polynomial IV specification, with county, state-month, month-year, and day-of-week fixed effects. The curve is concave, with the marginal effect increasing up to a turning point of approximately $23.7 \mu\text{g}/\text{m}^3$, above which it begins to decline.

Variable	Mean	Std. Dev.	Observations
Fatal crash	0.067	0.00013	4548823
$PM_{2.5}$ ($\mu g/m^3$)	10.60	7.14	4548823
tmax (C)	19.55	10.63	4548823
tmin (C)	7.18	9.84	4548823
Wind Speed (m/s)	4.49	2.87	4548823
Precipitation	4.50	2.87	4548823

Table 1.1: Summary Statistics: Daily Analysis, 1999–2013

Notes: tmax and tmin are reported here as continuous variables for descriptive purposes. In the regression analysis, they are entered as binned indicators (17 bins each, spanning -15°C to 30°C) interacted with precipitation and wind speed decile indicators.

Variable	Mean	Std. Dev.	Observations
Fatal crash	0.00358	0.06013	5091590
$PM_{2.5}$ ($\mu g/m^3$)	9.60100	7.14000	5091590
tmax (C)	19.45300	10.87249	5091590
tmin (C)	7.11070	9.97050	5091590
Wind Speed (m/s)	4.58600	2.87000	5091590
Precipitation	4.29600	2.87000	5091590

Table 1.2: Summary Statistics: Hourly Analysis, 2010–2013

Notes: tmax and tmin are reported here as continuous variables for descriptive purposes. In the regression analysis, they are entered as binned indicators (17 bins each, spanning -15°C to 30°C) interacted with precipitation and wind speed decile indicators.

Table 1.3: OLS Regression Results using Daily Data

	Dependent Variable: Total Car Crashes	
	(1)	(2)
$PM_{2.5}$	0.0002*** (0.00003)	0.0001*** (0.00003)
Observations	4548815	4548815
Controls	×	✓
Fixed Effects	×	✓

Note: Results are based on daily-level data for the US. The dependent variable is the total number of fatal car crashes per day. Column (1) presents the baseline OLS regression without controls or fixed effects. Column (2) includes additional weather-related controls and fixed effects. Standard errors are clustered at the county level, and all regressions are weighted by the county population. Significance levels are denoted by * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table 1.4: OLS Regression Results using hourly data

	Dependent Variable: Fatal Car Crashes	
	(1)	(2)
PM _{2.5}	0.000018*** (0.000003)	0.000009*** (0.000003)
Observations	5,091,589	5,091,589
Residual Std. Error	0.0603 (df = 5,091,588)	0.0603 (df = 5,091,588)
F Statistic	37.00*** (df = 1; 5,091,588)	52.30*** (df = 2; 5,091,588)
Controls	✗	✓
Fixed Effects	✗	✓

Note: Results in this table are based on hourly-level data for the US for the period 2010-2013. The dependent variable is the number of fatal car crashes per hour. Column (1) presents the baseline OLS regression without controls or fixed effects. Column (2) includes additional weather-related controls and fixed effects. Standard errors are clustered at the county level, and all regressions are weighted by the county population. Significance levels are denoted by *p < 0.1; **p < 0.05; ***p < 0.01.

Table 1.5: The Effect of $PM_{2.5}$ on Fatal Car Crashes: IV Estimates at Daily Level

	Dependent Variable: Fatal Car Crashes		
	(1)	(2)	(3)
$PM_{2.5}$	0.003639*** (0.000189)	0.000895*** (0.000023)	0.000905*** (0.000023)
Precipitation		-0.00260*** (0.00005)	-0.00260*** (0.00005)
Wind Speed		-0.00006 (0.00005)	-0.00005 (0.00005)
County FE	✗	✓	✓
State-Month FE	✗	✓	✓
Month-Year FE	✗	✓	✓
Day-of-Week FE	✗	✗	✓
Mean of Crashes	0.067	0.067	0.067
% Effect	5.43%	1.34%	1.35%
Observations	4,548,815	4,548,815	4,548,814
Residual Std. Error	0.3305	0.2854	0.2852
F Statistic	370.5***	2,940***	3,044***

Note: Results in this table are based on daily-level data for the U.S. The dependent variable is the number of fatal car crashes per day. Column (1) presents the baseline IV regression without controls or fixed effects. Column (2) adds controls for weather conditions and temperature, along with County, State-Month, and Month-Year fixed effects. Column (3) adds Day-of-Week fixed effects to account for variations in driving patterns across days. Each coefficient represents the effect of a $1 \mu\text{g}/\text{m}^3$ increase in $PM_{2.5}$ on fatal crashes. The percentage effects are calculated relative to the mean number of crashes (0.067 per day). Standard errors are clustered at the county level. In all columns, the instrument for $PM_{2.5}$ is constructed as the interaction of wind direction bins with spatial monitor group indicators (k-means clusters), so group-level fixed effects are implicitly absorbed by the instrument in all specifications. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.6: The Effect of $PM_{2.5}$ on Fatal Car Crashes: IV Estimates at Daily Level, 2010–2013

	Dependent Variable: Fatal Car Crashes		
	(1)	(2)	(3)
$PM_{2.5}$	0.003954*** (0.000406)	0.000694 (0.000368)	0.000734* (0.000367)
Precipitation		−0.000120 (0.000065)	−0.000078 (0.000065)
Wind Speed		−0.000075 (0.000282)	−0.000000 (0.000281)
County FE	✗	✓	✓
State-Month FE	✗	✓	✓
Month-Year FE	✗	✓	✓
Day-of-Week FE	✗	✗	✓
Mean of Crashes	0.067	0.067	0.067
% Effect	5.90%	1.04%	1.10%
Observations	855,055	855,055	855,055
Residual Std. Error	0.2987	0.2754	0.2751
F Statistic	5,208.6***	7,632.9***	7,698.7***

Note: Results in this table are based on daily-level data for the U.S. for the period 2010–2013, for direct comparison with the hourly analysis. The dependent variable is the number of fatal car crashes per day. Column (1) presents the baseline IV regression without controls or fixed effects. Column (2) adds controls for weather conditions and temperature, along with County, State-Month, and Month-Year fixed effects. Column (3) adds Day-of-Week fixed effects. Each coefficient represents the effect of a $1 \mu\text{g}/\text{m}^3$ increase in $PM_{2.5}$ on fatal crashes. The percentage effects are calculated relative to the mean number of crashes (0.067 per day). Standard errors are clustered at the county level. In all columns, the instrument for $PM_{2.5}$ is constructed as the interaction of wind direction bins with spatial monitor group indicators (k-means clusters), so group-level fixed effects are implicitly absorbed by the instrument in all specifications. Significance levels: $p < 0.1$, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Table 1.7: The Effect of PM_{2.5} on Fatal Car Crashes: IV Estimates at Daily Level, 2010–2013 (8 Wind Direction Bins)

	Dependent Variable: Fatal Car Crashes		
	(1)	(2)	(3)
PM _{2.5}	0.004451*** (0.000386)	0.000767* (0.000345)	0.000847* (0.000343)
Precipitation		−0.000111 (0.000063)	−0.000065 (0.000063)
Wind Speed		−0.000032 (0.000271)	0.000067 (0.000270)
County FE	✗	✓	✓
State-Month FE	✗	✓	✓
Month-Year FE	✗	✓	✓
Day-of-Week FE	✗	✗	✓
Mean of Crashes	0.067	0.067	0.067
% Effect	6.64%	1.14%	1.26%
Observations	855,055	855,055	855,055
Residual Std. Error	0.2989	0.2754	0.2751
F Statistic	2,479.3***	3,737.5***	3,775.5***

Note: Results in this table are based on daily-level data for the U.S. for the period 2010–2013, using 8 wind direction bins as instruments. The dependent variable is the number of fatal car crashes per day. Column (1) presents the baseline IV regression without controls or fixed effects. Column (2) adds controls for weather conditions and temperature, along with County, State-Month, and Month-Year fixed effects. Column (3) adds Day-of-Week fixed effects. Each coefficient represents the effect of a 1 $\mu\text{g}/\text{m}^3$ increase in PM_{2.5} on fatal crashes. The percentage effects are calculated relative to the mean number of crashes (0.067 per day). Standard errors are clustered at the county level. In all columns, the instrument for PM_{2.5} is constructed as the interaction of wind direction bins with spatial monitor group indicators (k-means clusters), so group-level fixed effects are implicitly absorbed by the instrument in all specifications. Significance levels: $\cdot p < 0.1$, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Table 1.8: Heterogeneity by Pollution Level: IV Estimates, Low PM_{2.5} Sample (Below Median)

	Dependent Variable: Fatal Car Crashes		
	(1)	(2)	(3)
PM _{2.5}	0.04502 (0.001378)	0.001659 (0.001382)	0.001694 (0.001378)
Precipitation		−0.000061 (0.000058)	−0.000047 (0.000058)
Wind Speed		0.000011 (0.000304)	0.000045 (0.000303)
County FE	✗	✓	✓
State-Month FE	✗	✓	✓
Month-Year FE	✗	✓	✓
Day-of-Week FE	✗	✗	✓
Mean of Crashes	0.067	0.067	0.067
% Effect	—	2.48%	2.53%
Observations	1,262,413	1,262,412	1,262,412
Residual Std. Error	0.3055	0.2756	0.2754
F Statistic	4,159.3***	4,662.4***	4,683.2***

Note: Sample restricted to county-days with PM_{2.5} below the median of 8.92 $\mu\text{g}/\text{m}^3$. Column (1) presents the baseline IV regression without controls or fixed effects. Columns (2) and (3) add weather controls and fixed effects. Standard errors are clustered at the county level. In all columns, the instrument for $PM_{2.5}$ is constructed as the interaction of wind direction bins with spatial monitor group indicators (k-means clusters), so group-level fixed effects are implicitly absorbed by the instrument in all specifications. Significance levels: $\cdot p < 0.1$, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Table 1.9: Heterogeneity by Pollution Level: IV Estimates, High PM_{2.5} Sample (Above Median)

	Dependent Variable: Fatal Car Crashes		
	(1)	(2)	(3)
PM _{2.5}	−0.001291 (0.000550)	0.000747* (0.000365)	0.000825* (0.000364)
Precipitation		−0.000090 (0.000084)	−0.000039 (0.000083)
Wind Speed		−0.000182 (0.000250)	−0.000050 (0.000250)
County FE	✗	✓	✓
State-Month FE	✗	✓	✓
Month-Year FE	✗	✓	✓
Day-of-Week FE	✗	✗	✓
Mean of Crashes	0.067	0.067	0.067
% Effect	—	1.11%	1.23%
Observations	1,262,412	1,262,411	1,262,411
Residual Std. Error	0.3607	0.3251	0.3247
F Statistic	2,859.5***	6,510.6***	6,527.2***

Note: Sample restricted to county-days with PM_{2.5} above the median of 8.92 $\mu\text{g}/\text{m}^3$. Column (1) presents the baseline IV regression without controls or fixed effects. Columns (2) and (3) add weather controls and fixed effects. Standard errors are clustered at the county level. In all columns, the instrument for $PM_{2.5}$ is constructed as the interaction of wind direction bins with spatial monitor group indicators (k-means clusters), so group-level fixed effects are implicitly absorbed by the instrument in all specifications. Significance levels: $\cdot p < 0.1$, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Table 1.10: Heterogeneity by Pollution Level: IV Estimates by Tercile

	Dependent Variable: Fatal Car Crashes		
	T1 (0–6.76)	T2 (6.76–11.7)	T3 (11.7+)
PM _{2.5}	0.002933 (0.003242)	0.005876* (0.002804)	0.000341 (0.000573)
Precipitation	–0.000016 (0.000081)	–0.000055 (0.000074)	–0.000179 (0.000124)
Wind Speed	0.000192 (0.000452)	–0.000178 (0.000295)	–0.000111 (0.000321)
County FE	✓	✓	✓
State-Month FE	✓	✓	✓
Month-Year FE	✓	✓	✓
Day-of-Week FE	✓	✓	✓
Mean of Crashes	0.067	0.067	0.067
% Effect	4.38%	8.77%	0.51%
Observations	841,608	841,608	841,605
Residual Std. Error	0.2587	0.3113	0.3284
F Statistic	1,305.7***	2,213.3***	2,480.2***

Note: Sample split into terciles of $PM_{2.5}$ concentration. Tercile cutpoints are 6.76 and 11.7 $\mu\text{g}/\text{m}^3$. All specifications include weather controls, County, State-Month, Month-Year, and Day-of-Week fixed effects. Standard errors are clustered at the county level. In all columns, the instrument for $PM_{2.5}$ is constructed as the interaction of wind direction bins with spatial monitor group indicators (k-means clusters), so group-level fixed effects are implicitly absorbed by the instrument in all specifications. Significance levels: $\cdot p < 0.1$, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Table 1.11: Heterogeneity by Pollution Level: IV Estimates by Quartile

	Dependent Variable: Fatal Car Crashes			
	Q1 (0–5.76)	Q2 (5.76–8.92)	Q3 (8.92–13.6)	Q4 (13.6+)
PM _{2.5}	0.003936 (0.005842)	0.003460 (0.006910)	0.011797** (0.004413)	0.000098 (0.000789)
Precipitation	–0.000038 (0.000107)	–0.000071 (0.000081)	0.000037 (0.000097)	–0.000237 (0.000162)
Wind Speed	0.000277 (0.000617)	–0.000250 (0.000366)	–0.000000 (0.000351)	–0.000070 (0.000388)
County FE	✓	✓	✓	✓
State-Month FE	✓	✓	✓	✓
Month-Year FE	✓	✓	✓	✓
Day-of-Week FE	✓	✓	✓	✓
Mean of Crashes	0.067	0.067	0.067	0.067
% Effect	5.87%	5.16%	17.61%	0.15%
Observations	631,206	631,205	631,205	631,202
Residual Std. Error	0.2459	0.3017	0.3195	0.3298
F Statistic	497.4***	802.7***	1,029.5***	1,237.3***

Note: Sample split into quartiles of $PM_{2.5}$ concentration. Quartile cutpoints are 5.76, 8.92, and 13.6 $\mu\text{g}/\text{m}^3$. All specifications include weather controls, County, State-Month, Month-Year, and Day-of-Week fixed effects. Standard errors are clustered at the county level. In all columns, the instrument for $PM_{2.5}$ is constructed as the interaction of wind direction bins with spatial monitor group indicators (k-means clusters), so group-level fixed effects are implicitly absorbed by the instrument in all specifications. Significance levels: $\cdot p < 0.1$, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Table 1.12: Polynomial IV: Effect of $PM_{2.5}$ on Fatal Car Crashes (Full Sample)

	(1)	(2)	(3)
$PM_{2.5}$	0.006888 (0.004856)	0.000889 (0.000471)	0.000980* (0.000474)
$PM_{2.5}^2$	-0.000122 (0.000117)	-0.000019 (0.000011)	-0.000021 (0.000011)
Precipitation		-0.000098* (0.000045)	-0.000072 (0.000045)
Wind Speed		-0.000192 (0.000177)	-0.000109 (0.000177)
County FE	X	✓	✓
State-Month FE	X	✓	✓
Month-Year FE	X	✓	✓
Day-of-Week FE	X	X	✓
Observations	2,524,825	2,524,825	2,524,825
F-test (1st stage, $PM_{2.5}$)	2,403,965***	1,712,568***	1,710,852***
F-test (1st stage, $PM_{2.5}^2$)	547,862***	466,249***	466,047***

Note: Results are based on daily-level data for the U.S. The dependent variable is the number of fatal car crashes per day. $PM_{2.5}$ and $PM_{2.5}^2$ are instrumented using four cardinal wind direction bins (N, E, S, W) and their interactions with $PM_{2.5}$ concentration, yielding six instruments for two endogenous variables. Standard errors clustered at the county level are reported in parentheses. Significance levels: $p < 0.1$, $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Table 1.13: The Effect of PM_{2.5} on Fatal Car Crashes: IV Estimates at hourly level

	Dependent Variable: Fatal Car Crashes		
	(1)	(2)	(3)
PM _{2.5}	0.000265*** (0.000034)	0.000016* (0.000010)	0.000071*** (0.000009)
Precipitation		-0.0002*** (0.00001)	-0.000152*** (0.000009)
Wind Speed		-0.0001*** (0.00001)	-0.000081*** (0.00001)
County FE	✗	✓	✓
State-Month FE	✗	✓	✓
Month-Year FE	✗	✓	✓
Day-of-Week FE	✗	✗	✓
Mean of Crashes	0.0036	0.0036	0.0036
% Effect	7.36%	0.44%	1.97%
Observations	5,091,589	5,091,589	5,091,589
Residual Std. Error	0.0603	0.0603	0.06026
F Statistic	59.91***	192.50***	374.3***

Note: Results in this table are using hourly-level data for the US for the period 2010-2013. The dependent variable is the number of fatal car crashes per hour. Column (1) presents the baseline IV regression without any controls or fixed effects. Column (2) includes additional controls for weather conditions and temperature, along with County, State-Month, and Month-Year FE. Column (3) adds Day-of-Week fixed effects to account for variation in driving patterns across different days of the week. Each coefficient represents the effect of a 1 $\mu\text{g}/\text{m}^3$ increase in PM_{2.5} on fatal crashes. The percentage effects are calculated relative to the mean number of crashes (0.0036 per hour). Standard errors are clustered at the county level, and all regressions are weighted by the county population. In all columns, the instrument for PM_{2.5} is constructed as the interaction of wind direction bins with spatial monitor group indicators (k-means clusters), so group-level fixed effects are implicitly absorbed by the instrument in all specifications. Significance levels are denoted by *p < 0.1, **p < 0.05, ***p < 0.01*.

Table 1.14: Time-of-Day Heterogeneity: Effect of PM_{2.5} on Fatal Car Crashes

	(1) Morning Rush <i>7am–9am</i>	(2) Midday <i>10am–3pm</i>	(3) Evening Rush <i>4pm–7pm</i>	(4) Late Night <i>8pm–6am</i>
PM _{2.5}	$6.67 \times 10^{-5}\dagger$ (3.91×10^{-5})	1.53×10^{-5} (3.96×10^{-5})	$1.35 \times 10^{-4*}$ (5.55×10^{-5})	-1.61×10^{-6} (3.36×10^{-5})
County FE	Yes	Yes	Yes	Yes
Day-of-Week FE	Yes	Yes	Yes	Yes
State $\times$ Month FE	Yes	Yes	Yes	Yes
Month $\times$ Year FE	Yes	Yes	Yes	Yes
Mean of Crashes	0.0036	0.0036	0.0036	0.0036
% Effect	1.86%	0.43%	3.76%	−0.04%
Observations	632,829	1,263,941	851,667	2,342,673
Within R^2	5.11×10^{-6}	6.12×10^{-6}	2.74×10^{-5}	3.73×10^{-6}

Notes: Results are based on hourly-level data for the US for the period 2010–2013. The dependent variable is the number of fatal car crashes per hour. Each column presents IV estimates for a distinct time-of-day bin. Wind direction interacted with spatial pollution monitor clusters serves as an instrument for PM_{2.5}. All specifications include weather controls (wind speed, precipitation, maximum and minimum temperature). Standard errors are clustered at the county level. Significance levels: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, $\dagger p < 0.1$.

Table 1.15: Weekday vs. Weekend: Effect of PM_{2.5} on Fatal Car Crashes

	(1) Weekday	(2) Weekend
PM _{2.5}	$3.88 \times 10^{-5}\dagger$ (2.49×10^{-5})	$5.29 \times 10^{-5}\dagger$ (3.60×10^{-5})
County FE	Yes	Yes
Day-of-Week FE	Yes	Yes
State $\times$ Month FE	Yes	Yes
Month $\times$ Year FE	Yes	Yes
Mean of Crashes	0.0036	0.0036
% Effect	1.08%	1.48%
Observations	3,632,525	1,458,585
Within R^2	3.66×10^{-6}	1.62×10^{-5}

Notes: Results are based on hourly-level data for the US for the period 2010–2013. The dependent variable is the number of fatal car crashes per hour. Column (1) restricts to Monday–Friday and Column (2) to Saturday–Sunday. Wind direction interacted with spatial pollution monitor clusters is as an instrument for PM_{2.5}. All specifications include weather controls (wind speed, precipitation, maximum and minimum temperature). Standard errors are clustered at the county level. Significance levels: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, $\dagger p < 0.1$.

Chapter 2:

Air Pollution and Traffic Crashes: Evidence from Chicago

Irina Firsova

University of California, Davis

Abstract

This paper examines the causal effect of fine particulate matter ($\text{PM}_{2.5}$) on traffic crashes in Chicago from 2017 to 2025. Using an instrumental variable approach that exploits exogenous variation in wind direction, I find that a $1 \mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ leads to a 0.78% increase in total crashes and a 9.05% increase in fatal crashes. The effect suggests that air pollution affects both crash frequency and severity. Individual-level analysis reveals that $\text{PM}_{2.5}$ increases crash severity conditional on a crash occurring, with higher pollution associated with more severe injuries.

2.1 Introduction

Traffic crashes impose substantial costs on society through loss of life, injuries, public health expenditures and economic damages. In the United States, motor vehicle crashes are a leading cause of death, with approximately 40,000 fatalities annually. According to the National Highway Traffic Safety Administration, the total economic cost of motor vehicle crashes was \$340 billion in 2019, rising to an estimated \$1.4 trillion when quality-of-life losses are included (Blincoe, Miller, Wang, Swedler, Coughlin, Lawrence, Guo, Klauer, and Dingus, 2023). These costs are comparable in scale to other major public health burdens: air pollution alone costs the U.S. economy an estimated \$790 billion annually in economic welfare losses (World Bank, 2016). Extensive research has examined how policy interventions affect traffic safety, demonstrating the effectiveness of mandatory seat-belt laws (Cohen and Einav, 2003) and drunk driving regulations (Eisenberg, 2003). However, less attention has been paid to how environmental factors influence crash risk. This paper contributes to a growing literature on the cognitive effects of air pollution by examining whether short-term exposure to fine particulate matter ($\text{PM}_{2.5}$) causally increases traffic crashes in Chicago.

This paper provides new evidence on the causal relationship between air pollution and traffic crashes by focusing on Chicago. It is a major metropolitan area with variation in both air quality and traffic patterns. Chicago provides an ideal setting for several reasons. First, the city has comprehensive crash data covering all reported incidents (both fatal and non-fatal) from 2017 through 2025. Prior studies, including Burton and Roach (2026), rely on the Fatality Analysis Reporting System (FARS), a national database that records only

fatal crashes. Similarly, [Sager \(2019\)](#) use data on vehicles involved in injury accidents, which excludes property-damage-only crashes. By contrast, the Chicago crash data includes every reported incident regardless of severity, enabling separate estimation of pollution’s effect on crash frequency and on the conditional probability of a fatality. It includes detailed individual-level information on crash severity, timing, location, and environmental conditions. Second, Chicago’s geographic position creates variation in pollution exposure driven by wind direction. The city is bounded by Lake Michigan to the east and northeast, while industrial facilities and major highways are concentrated to the south and southwest. This geographic setting means that winds from different directions bring different pollution levels, providing a source of exogenous variation for identification. Third, I measure pollution using ground-level EPA monitors rather than satellite-derived data. While satellite data offers broad geographic coverage, ground-based sensors provide more accurate surface-level measurements, which is what matters for drivers’ actual exposure. Within a single dense urban area like Chicago, the monitors provide reliable coverage of city-wide daily variation in $PM_{2.5}$. Averaging across monitors reduces noise from any individual station. Studies that aggregate data across many counties or states must face greater heterogeneity in monitor placement and local pollution sources in exposure, all of which introduce measurement error that attenuates estimated effects toward zero.

The main empirical challenge in estimating the causal effect of air pollution on crashes is that pollution exposure is not randomly assigned. Air pollution levels are correlated with economic activity, traffic volume, and weather conditions that independently affect crash risk, creating omitted variable bias in OLS estimates. The direction of this bias is ambiguous: higher traffic volume raises both pollution and crashes, pushing OLS upward,

while weather conditions that clear pollution (such as rain and wind) may simultaneously increase crash risk, pushing OLS downward. At the same time, pollution monitors may be strategically placed (Sullivan et al., 2018), and measured $\text{PM}_{2.5}$ at fixed monitor locations may not represent true exposure for drivers distributed across the city, attenuating OLS estimates toward zero. To address these concerns, I employ an instrumental variables (IV) approach that exploits exogenous variation in daily wind direction. Following Deryugina et al. (2019), I use wind direction as an instrument for $\text{PM}_{2.5}$ concentrations. The key idea is that wind direction affects local pollution concentrations by transporting pollutants from different source areas, but wind direction itself does not directly affect driving behavior or crash risk once I control for wind speed and other weather conditions (Knapp et al., 2000).

The analysis proceeds in two stages. I examine both whether crashes occur and how severe they are, where severity is measured using the five ordered injury categories assigned by responding officers: No Injury, Reported Not Evident, Non-Incapacitating, Incapacitating, and Fatal. First, I examine the determinants of crash severity using individual-level crash data. The Chicago crash records include detailed information on posted speed limits, lighting conditions, road surface, and weather at the time of each crash. These granular controls are not available in national fatal-crash databases. This analysis identifies which of these factors contribute to more severe injuries and provides context for interpreting the pollution effects. I find that posted speed limits, darkness, certain weather conditions, and off-peak timing all increase crash severity. Importantly, I also find evidence that $\text{PM}_{2.5}$ increases crash severity conditional on a crash occurring: a $1 \text{ } \mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ raises the odds of being in a more severe injury category by 0.3%. While this effect is smaller than the effects of other environmental factors (for comparison, crashes on dark but lighted roads have 37% higher

odds of severe injury and crashes on wet surfaces have 19% higher odds), it provides direct evidence that pollution impairs the cognitive functions necessary for crash mitigation.

Second, I estimate the causal effect of $\text{PM}_{2.5}$ on the number of crashes using daily aggregated data and an IV approach that uses wind direction bins as instruments for pollution. I find that a $1 \text{ } \mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ leads to 2.3 additional total crashes per day (a 0.78% increase relative to the daily mean of 292 crashes) and 0.028 additional fatal crashes per day (a 9.05% increase relative to the daily mean of 0.33 fatal crashes). This suggests that pollution affects not just crash frequency but also crash severity. This finding is consistent with the individual-level analysis showing that pollution increases the odds of more severe injuries conditional on a crash occurring. The first-stage F-statistic of 37.8 confirms that wind direction bins are strong predictors of $\text{PM}_{2.5}$ concentrations.

These results make several contributions to the literature. First, they provide causal evidence that air pollution significantly increases traffic fatalities, highlighting an important external cost of poor air quality. The traffic safety benefits of pollution reduction should be incorporated into cost-benefit analyses of environmental regulations. Using the Environmental Protection Agency’s value of a statistical life of \$7.4 million ([EPA, 2006](#)), a one standard deviation reduction in $\text{PM}_{2.5}$ in Chicago would prevent approximately 47 traffic fatalities per year. Monetary benefits would be \$348 million annually from reduced traffic deaths alone. Second, the paper contributes to the growing literature on pollution’s cognitive effects by examining a real-time decision-making context where impairment has immediate and potentially fatal consequences. Previous research has documented pollution’s effects on labor productivity ([He et al., 2019](#)), crime ([Herrnstadt et al., 2021](#)), and professional performance

([Archsmith et al., 2018](#)). This paper demonstrates that cognitive impairment leads to increased mortality through impaired driving. Third, by examining both individual-level crash severity and aggregate crash frequency, the paper shows the mechanisms linking pollution exposure to traffic fatalities. Pollution affects both whether crashes occur through impaired attention and how severe they are when they do occur through impaired reaction times that prevent evasive action.

The hypothesis that air pollution affects driving safety is motivated by the evidence documenting the adverse effects of pollution on human health and cognitive function. Air pollution is a major public health challenge. Multiple studies have shown that exposure to air pollution leads to adverse health outcomes through respiratory and cardiovascular diseases ([Kim et al., 2013](#); [Fiordelisi et al., 2017](#)). Long-term exposure to $PM_{2.5}$ has been linked to premature mortality, particularly among individuals with chronic heart or lung conditions ([Pope et al., 2002](#); [Brook et al., 2010](#); [Dockery et al., 1993](#); [Hoek et al., 2013](#); [Pope and Dockery, 2006](#)). These health impacts impose substantial economic costs through increased healthcare expenditures and reduced life expectancy ([Segalowitz, 2008](#)).

Beyond these well-documented health effects, recent research has revealed that air pollution significantly impairs cognitive function. Correlational studies have shown that pollution exposure is associated with reduced labor productivity through both increased absenteeism and worsened performance ([Hausman et al., 1984](#); [Carson et al., 2011](#); [Rodrigues-Silva, de Paula Santos, Saldiva, Amato-Lourenço, Miraglia et al., 2012](#)). More recent work using causal identification strategies reaches similar conclusions. [Hanna and Oliva \(2015\)](#) exploit a natural experiment in Mexico City, and [He et al. \(2019\)](#) use variation in indus-

trial activity in China, both finding that pollution reduces labor productivity. [Herrnstadt et al. \(2021\)](#) find that air pollution increases criminal activity in Chicago through its effects on cognitive function and self-control, using a wind direction instrument similar to the one employed here. [Archsmith et al. \(2018\)](#) document that pollution impairs decision-making among highly skilled professionals such as baseball umpires. The mechanism underlying these effects appears to work through pollution’s effect on cognitive processes such as attention and executive function ([Weuve et al., 2012](#); [Calderon-Garciduenas et al., 2008](#)).

Driving is a cognitively demanding task that requires sustained attention, rapid information processing, quick reflexes, and real-time decision-making under uncertainty. If air pollution impairs these cognitive functions even temporarily, it is reasonable to expect negative impacts on driving performance. The potential consequences are particularly severe because driving occurs in a high-stakes environment where even a small delay in reaction time can result in fatal crashes. The errors documented by [Archsmith et al. \(2018\)](#) among baseball umpires occur in relatively low-risk settings, but similar cognitive impairments among drivers could lead to more severe consequences.

The existing research on air pollution and mortality has predominantly focused on internal causes of death such as cardiovascular and respiratory disease, along with associated hospitalizations ([Deryugina et al., 2019](#)). However, pollution’s consequences extend beyond these direct health outcomes to include external causes of death such as traffic accidents. Two recent working papers examine this question: [Burton and Roach \(2026\)](#) analyze the relationship between $PM_{2.5}$ exposure and fatal crashes using national U.S. data. Also, [Danagoulian and Deza \(2024\)](#) examine how seasonal pollen affects traffic fatalities. The most closely re-

lated published study is by [Sager \(2019\)](#), who examine the impact of air pollution on road safety in the United Kingdom between 2009 and 2014. Using temperature inversions as an instrument for pollution, they find that a 1 $\mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ leads to a 0.3–0.6% increase in vehicles involved in road accidents per day.

A key limitation of prior work is that it focuses exclusively on fatal crashes, but fatal crashes is an extreme outcome. The number of fatal crashes can be written as:

$$N_{\text{fatal}} = N_{\text{total}} \times \Pr(\text{fatality} \mid \text{crash})$$

Prior studies estimate $\frac{dN_{\text{fatal}}}{d\text{Pollution}}$, but do not decompose this into $\frac{dN_{\text{total}}}{d\text{Pollution}}$ (the effect on crash frequency) and $\frac{d\Pr(\text{fatality}|\text{crash})}{d\text{Pollution}}$ (the effect on crash severity conditional on a crash occurring). This distinction is important for understanding the mechanism. If pollution primarily impairs attention and situational awareness, we would expect it to increase total crashes; if it impairs reaction time and the ability to take evasive action, we would expect it to increase severity conditional on a crash. Chicago’s comprehensive crash data, which records all incidents regardless of severity, allows me to estimate both margins separately.

The remainder of this paper proceeds as follows. Section 2 describes the data sources and provides descriptive statistics on crash patterns and air quality in Chicago. Section 3 presents the empirical strategy for estimating the causal effect of $\text{PM}_{2.5}$ on crash frequency and discusses the IV results. Section 4 examines the determinants of crash severity using individual-level data and places the pollution effects in the context of other environmental factors. Section 5 discusses the policy implications and avenues for future research.

2.2 Data

The analysis combines three data sources: traffic crash records from the Chicago Data Portal, $PM_{2.5}$ measurements from EPA air quality monitors, and meteorological data from weather stations. This section describes each source in turn and presents descriptive statistics.

2.2.1 Crash Data

Crash information is from the Chicago Data Portal (City of Chicago, Department of Transportation). The data include every reported motor vehicle crash within city limits from 2017 through 2025. Each observation has the date, time, location, severity classification, and contributing factors. Severity is classified by the responding officer into one of five ordered injury categories. Contributing factors recorded include distracted driving, failure to yield, improper lane use, and speeding, among others.

Table 1 presents summary statistics for the key variables. Columns report the number of observations, mean, standard deviation, minimum, 25th percentile, median, 75th percentile, and maximum for daily observations of traffic crashes, weather conditions, and air pollution levels in Chicago.

The outcome of interest is the number of total (fatal and non-fatal) crashes per day. The mean is 291.77, with a standard deviation of approximately 63.95. This shows a relatively high baseline number of daily traffic incidents, which is consistent with Chicago’s population, traffic, and infrastructure. The interquartile range is from 255 to 333, suggesting a relatively

skewed distribution with a long right tail, indicating occasional increases in crashes due to bad weather, traffic disruptions, or other external shocks.

Separating crashes into fatal and non-fatal incidents shows that most daily crashes are non-fatal. The mean number of fatal crashes is only 0.33, with a standard deviation of 0.60. The median is zero, and the 75th percentile is one, indicating that fatal crashes are rare and on most days they do not occur. This is consistent with expectations for urban settings, where low-speed collisions are more common and less likely to result in fatalities.

Table 2 presents summary statistics for crashes separated by injury severity. First, the vast majority of crashes result in no injury (85.8%), consistent with the majority of city collisions being low-speed. Fatal crashes are rare, representing only 0.11% of all incidents, or approximately 1,024 fatalities over the sample period. Other categories include reported but not evident injuries (4.5%), non-incapacitating injuries (7.9%), and incapacitating injuries (1.6%).

Table 3 provides additional context by showing the distribution of weather and lighting conditions across all crashes. The data shows that most crashes occur under favorable conditions: 78.4% happen in clear weather and 64.1% in daylight.

Figure 1 plots the total number of daily crashes in Chicago. The figure shows high daily variation. First, there is a clear upward trend in crashes during 2017 and 2018, which slows down around 2019. The period around early 2020 shows a sharp drop, consistent with the effects of the COVID-19 pandemic and the lockdown that led to a significant reduction in driving activity. After that, the number of daily crashes increases and stabilizes at a slightly

lower level than the pre-pandemic peak, with regular fluctuations around a mean of roughly 290 crashes per day. The data also shows seasonality and occasional spikes, which could be due to extreme weather, holidays, or other disruptions. While the overall trend is relatively stable post-2020, the graph suggests that the risk of crashes remains consistently high, with no evidence of long-term decline.

To put these numbers in context, Chicago is one of the most congested cities in the United States, with several major interstate highways, including I-90, I-94, I-290, and I-55, running through or around the city. High traffic volume means that a large number of crashes is expected even if the per-mile crash rate is not unusually high. The national fatality rate is approximately 1.20 deaths per 100 million vehicle miles traveled ([Blincoe et al., 2023](#)). The crash counts observed here reflect this high exposure: with millions of vehicle miles traveled daily in the Chicago metropolitan area, a mean of 292 total crashes per day corresponds to a crash rate that is broadly consistent with national urban benchmarks.

Figure 2 shows the average number of crashes per month. The monthly averages range from 270 in January to over 320 in September and October. There is a seasonal pattern, with lower average crash counts in winter months and an increase through spring and summer. Crash rates remain high in the fall and decline again in December. These patterns are consistent with seasonal changes in traffic volume and road conditions. The lower numbers in January and February likely reflect both reduced travel during cold months and cautious driving in icy or snowy conditions. In contrast, the summer and early fall months coincide with increased travel, more recreational driving, and more congestion.

Figure 3 shows the hourly distribution of crashes. Minimal incidents occur 02:00–05:00,

and counts significantly increase after 06:00 and reach the first peak around 08:00. A second peak occurs 16:00–19:00 during the evening commute. After 20:00 the frequency declines steadily. The variation in hourly crashes correlates with traffic flow.

Figure 4 shows the weekly distribution with minimum activity on Sunday and progressive increase toward Friday. Peak counts occur Thursday–Friday, followed by a small decline on Saturday. This reflects commuter travel concentration and increase in pre-weekend traffic volume.

2.2.2 Air Pollution Data

Air quality data are obtained from the U.S. Environmental Protection Agency’s Air Quality System (EPA, AQS). $PM_{2.5}$ refers to fine particulate matter with a diameter of 2.5 micrometers or less. It is the standard measure used in the health and economics literature because particles of this size penetrate deep into the lungs and can cross the blood-brain barrier, making them the most relevant pollutant for both health and cognitive outcomes ([Brook et al., 2010](#)). It is also the EPA’s primary regulatory standard for air quality. The analysis uses daily average $PM_{2.5}$ concentrations as the time unit, which is consistent with the medical literature showing that acute cognitive impairment from pollution exposure can occur within hours of exposure and persist over the course of a day ([Calderon-Garciduenas et al., 2008](#)). I use hourly $PM_{2.5}$ measurements from all monitors within Cook County and the surrounding collar counties. EPA monitors are sited according to federal regulatory guidelines designed to measure population exposure, rather than placed in response to local crash patterns, supporting their use as an exogenous measure of air quality. The collar counties (DuPage, Kane,

Lake, McHenry, and Will) surround Cook County and together with Cook County comprise the six-county Chicago metropolitan area. While crash data are only available for the city of Chicago, monitors from the collar counties are included because pollution from surrounding areas contributes to air quality within the city. $PM_{2.5}$ monitors are spatially matched to the city of Chicago by selecting monitors located within Cook County or within 20 km of the city border. Following [Deryugina et al. \(2019\)](#), I construct daily $PM_{2.5}$ concentrations by averaging hourly monitor readings across all available monitors on a given day. Averaging across monitors within Cook County and the collar counties reduces noise from any individual station and provides a more reliable measure of the pollution that Chicago drivers are exposed to. During the sample period, approximately 11 EPA monitors were active in Cook County and the surrounding collar counties. The measure of air quality used in the paper is the average $PM_{2.5}$ concentration in micrograms per cubic meter ($\mu\text{g}/\text{m}^3$). For descriptive purposes, I also compute the corresponding Air Quality Index (AQI).

As for environmental pollution indicators, the dataset includes daily averages for $PM_{2.5}$ concentrations and the Air Quality Index (AQI). The average $PM_{2.5}$ concentration is $9.18 \mu\text{g}/\text{m}^3$, with a standard deviation of 4.66. There was a spike in June 2023 caused by wildfires in Canada. The AQI averages 44.95. These levels suggest that air quality in Chicago typically falls within the “Good” to “Moderate” range.

Figure 5 shows the average $PM_{2.5}$ concentration by month in Chicago. The highest pollution levels occur in December and June, with monthly averages around $10 \mu\text{g}/\text{m}^3$. The lowest concentrations are recorded in September and October, falling below $9 \mu\text{g}/\text{m}^3$.

Figure 6 shows daily $PM_{2.5}$ concentrations in Chicago over time. The data show high

daily variation, with levels generally ranging between 5 and 15 $\mu\text{g}/\text{m}^3$. There are multiple short spikes, some reaching concentrations above 30 $\mu\text{g}/\text{m}^3$, and one extreme outlier in 2023 that exceeds 80 $\mu\text{g}/\text{m}^3$. This spike aligns with the timing of the Canadian wildfires, which affected large parts of the Midwest in summer 2023. The plot also shows that the average pollution levels remained fairly stable over time, with no clear long-term decline. To provide context, the EPA classifies daily $PM_{2.5}$ concentrations below 12 $\mu\text{g}/\text{m}^3$ as “Good,” 12.1–35.4 $\mu\text{g}/\text{m}^3$ as “Moderate,” 35.5–55.4 $\mu\text{g}/\text{m}^3$ as “Unhealthy for Sensitive Groups,” and above 55.5 $\mu\text{g}/\text{m}^3$ as “Unhealthy” for the general population (EPA, 2006). Most days in Chicago fall in the Good to Moderate range, but the spikes visible in Figure 6, including the 2023 wildfire event exceeding 80 $\mu\text{g}/\text{m}^3$, reach levels classified as Unhealthy.

2.2.3 Weather Data

Weather data are obtained from the City of Chicago Data Portal and NOAA. Additional weather observations are obtained from NOAA weather stations in the Chicago area. When multiple stations are available, observations are averaged to obtain a single representative daily value for each variable.

Weather related variables also show variation. Most days see little or no precipitation with occasional extreme events. The distribution is right-skewed, indicating that a very small number of days experienced very high precipitation.

The average daily wind direction is approximately 177° , pointing to a predominantly southern wind. The standard deviation is 70° , reflecting considerable variation. Average

wind speed is 2.72 m/s, ranging from 0.58 to 44.83 m/s.

2.3 Effects of Pollution on Crashes

The main empirical challenge in estimating the causal effect of air pollution on traffic crashes is that exposure to PM2.5 is not randomly assigned. Air pollution levels are correlated with various economic activities, traffic patterns, and weather conditions that independently affect crash risk. Moreover, pollution monitors may be strategically placed, and measured pollution at monitors may not represent true exposure for all drivers in the city. These issues introduce both measurement error and omitted variable bias in ordinary least squares (OLS) estimates. To address these concerns, I use an instrumental variables (IV) approach that exploits exogenous variation in daily wind direction. The key insight is that while wind direction affects local pollution concentrations by transporting pollutants from different source areas, wind direction itself does not directly affect driving behavior or crash risk when I control for wind speed and other weather conditions.

Let d index days and m denote month and y - year respectively. The relationship between crashes and fine particulate matter in Chicago is given by:

$$Y_{dmy} = \beta PM2.5_{dmy} + X'_{dmy}\gamma + \delta_w + \theta_{my} + \epsilon_{dmy}$$

The specification includes day-of-week fixed effects to account for systematic differences in traffic patterns and driving behavior across weekdays and weekends, and month-by-year

fixed effects to control for seasonal variation in both crashes and pollution, as well as year-to-year trends. These month-by-year fixed effects are particularly important given that the sample excludes 2020 and 2021 due to the COVID-19 pandemic’s effects on traffic patterns. With daily aggregation, I cannot include day fixed effects, but control for temporal patterns through month-by-year and day-of-week fixed effects. The vector X'_{dmy} includes the following weather controls: precipitation intensity (mm/hr), total daily precipitation (mm), wind speed (m/s), and average daily temperature ($^{\circ}\text{C}$).

The use of daily data is appropriate because there is substantial day-to-day variation in both $\text{PM}_{2.5}$ and crash counts within month-by-year and day-of-week cells. The overall standard deviation of daily $\text{PM}_{2.5}$ is $4.66 \mu\text{g}/\text{m}^3$ around a mean of $9.18 \mu\text{g}/\text{m}^3$, reflecting large changes in air quality from one day to the next that are not fully explained by seasonal patterns or day of week. Similarly, daily crash counts have a standard deviation of 63.95 around a mean of 292. Month-by-year and day-of-week fixed effects absorb seasonal and weekly patterns, but considerable idiosyncratic variation is left, providing the identifying variation used in estimation.

The instrument is constructed from daily average wind direction, measured in degrees (0–360 $^{\circ}$). As reported in Table 1, wind direction has a mean of 177 $^{\circ}$ and a standard deviation of 70 $^{\circ}$, reflecting substantial day-to-day variation in the direction from which air arrives in Chicago. To convert this continuous measure into instruments, I divide the compass into discrete bins. The main specification uses four 90 $^{\circ}$ bins: North [0 $^{\circ}$ –90 $^{\circ}$), East [90 $^{\circ}$ –180 $^{\circ}$), South [180 $^{\circ}$ –270 $^{\circ}$), and West [270 $^{\circ}$ –360 $^{\circ}$). A robustness specification uses eight 45 $^{\circ}$ bins, providing finer geographic resolution at the cost of additional instruments and a weaker first

stage.

The identifying assumption for the IV approach is that conditional on weather controls and fixed effects, daily wind direction affects crashes only through its impact on air pollution. This assumption is plausible because wind direction determines which pollution sources contribute to local concentrations, but wind direction itself, holding wind speed constant, does not directly affect driving conditions or behavior. I construct instruments by dividing wind direction into discrete bins. Following the approach in Deryugina et al. (2019), I create indicator variables for wind direction bins, with one bin serving as the omitted reference category. While Deryugina et al. (2019) apply this approach at the county-group level aggregating across many counties, the analysis here operates at the city level using daily observations for Chicago. This difference is appropriate given the research question. Studying a single city avoids the cross-county heterogeneity in pollution sources, traffic patterns, and demographic composition that complicates interpretation in multi-county analyses. Wind direction generates meaningful variation in $PM_{2.5}$ within Chicago because the city’s pollution sources are geographically concentrated.

$$PM2.5_{dmy} = \sum_{b \in B} \alpha_b \mathbb{1}[WindDir_{dmy} = b] + X'_{dmy} \pi + \delta_w + \theta_{my} + u_{dmy}$$

This specification uses three instruments (four bins minus the reference category).

This IV strategy requires that wind direction bins are strong predictors of pollution levels and that they satisfy the exclusion restriction. The first-stage F-statistic tests instrument strength.

2.3.1 First Stage Results

Chicago's pollution sources are not uniformly distributed, which provides the geographic variation exploited by the IV strategy. Industrial facilities and major highways are concentrated to the south and southwest of the city center, while Lake Michigan lies to the east and northeast. Consequently, southerly and southwesterly winds tend to bring higher PM2.5 concentrations, while easterly and northerly winds from the lake bring cleaner air.

Figure 7 shows the first stage relationship between wind direction and PM2.5 concentrations. The pattern is clear: winds from the east (E) and southeast (SE) are associated with PM2.5 concentrations above 10 $\mu\text{g}/\text{m}^3$, while winds from the northwest (NW) are associated with concentrations around 5 $\mu\text{g}/\text{m}^3$. This substantial variation across wind directions provides strong predictive power for the instrumental variables approach.

2.3.2 OLS Results

Table 5 presents OLS estimates of the effect of PM2.5 on daily crashes. The OLS coefficient on PM2.5 is negative but statistically insignificant across all three outcomes (total crashes, fatal crashes, and non-fatal crashes). The weather controls show expected patterns: precipitation intensity is strongly positively associated with crashes (10.1 additional total crashes per unit increase in rain intensity), consistent with rain creating hazardous driving conditions. Wind speed is negatively associated with total and non-fatal crashes (2.2 fewer crashes per m/s increase in wind speed), potentially reflecting reduced travel during high-wind conditions. Wind speed also shows a negative relationship with fatal crashes,

though the magnitude is much smaller (0.020 fewer fatal crashes per m/s). The statistically insignificant OLS coefficients on PM2.5 suggest that OLS fails to capture the true effect of air pollution on crashes. This is consistent with measurement error attenuating OLS estimates toward zero, as pollution monitors may not accurately capture exposure for all drivers throughout the city.

2.3.3 IV Results

Table 6 presents the main IV results using four 90° wind direction bins as instruments. The IV estimates are dramatically different from OLS, both in magnitude and statistical significance. A 1 $\mu\text{g}/\text{m}^3$ increase in PM2.5 leads to 2.314 additional total crashes per day ($p < 0.05$) and 0.028 additional fatal crashes per day ($p < 0.05$). The effect on non-fatal crashes is nearly identical to total crashes (2.300 additional crashes), as expected given that fatal crashes are a small fraction of total crashes.

Relative to the sample means, these effects represent a 0.78% increase in total crashes and a 9.05% increase in fatal crashes per 1 $\mu\text{g}/\text{m}^3$ increase in PM2.5. The much larger percentage effect for fatal crashes suggests that air pollution has a disproportionate impact on crash severity, not just crash frequency.

The first-stage F-statistic is 37.8, exceeding the Stock-Yogo critical value of 22.3 for 10% maximal IV size with three instruments, indicating that wind direction bins are strong predictors of PM2.5 concentrations and that the IV estimates are reliable.

Given Chicago’s mean PM2.5 concentration of 9.18 $\mu\text{g}/\text{m}^3$ with a standard deviation of

4.66 $\mu\text{g}/\text{m}^3$, a one standard deviation increase in pollution would lead to approximately 10.8 additional total crashes and 0.13 additional fatal crashes per day. Over a year, this translates to roughly 47 additional fatalities attributable to one standard deviation higher pollution, a substantial public health burden.

Table 7 presents results using eight 45° wind direction bins as a robustness check. The 8-bin specification yields qualitatively similar but slightly smaller estimates: a 1 $\mu\text{g}/\text{m}^3$ increase in PM_{2.5} leads to 1.600 additional total crashes ($p < 0.10$) and 0.020 additional fatal crashes ($p < 0.10$). These translate to percentage effects of 0.54% for total crashes and 6.46% for fatal crashes. However, the 8-bin specification produces a weaker first-stage F-statistic of 17.6, which falls below the Stock-Yogo critical value of 22.3. This reduction in instrument strength adds both noise and potential bias to the estimates. Bias in 2SLS is proportional to the number of excluded instruments minus two, so the additional instruments in the 8-bin specification may introduce more bias than they reduce. For this reason, I focus on the 4-bin specification as my main result, though the consistency between the two specifications in terms of sign, magnitude, and the pattern of larger effects on fatal crashes provides reassurance that the findings are not driven by the specific choice of instruments.

The first chapter of this dissertation finds that a 1 $\mu\text{g}/\text{m}^3$ increase in PM_{2.5} leads to a 1.34% increase in daily fatal crashes nationally, smaller than the 9.05% effect on fatal crashes found here, which reflects the difference in setting. The national estimate pools urban and rural areas, while this paper focuses on a dense urban environment where cognitive impairment from pollution may have larger consequences due to complex driving conditions and higher traffic volume. [Sager \(2019\)](#) finds a 2% increase in total accidents per grid cell

per day for a 1 $\mu\text{g}/\text{m}^3$ increase in NO_2 in the UK, consistent with the 0.78% effect on total crashes found here. Direct comparison is complicated by the difference in pollutant, instrument (temperature inversions vs wind direction), and geographic unit. A key advantage of the present paper is the ability to separately identify effects on crash frequency and crash severity.

2.4 Effects of Pollution on Crash Severity

Before examining crash severity in depth, it is useful to recall the decomposition introduced in Section 1: $N_{\text{fatal}} = N_{\text{total}} \times \Pr(\text{fatality} \mid \text{crash})$. The IV results in Section 3 establish that pollution increases N_{total} . This section examines $\Pr(\text{fatality} \mid \text{crash})$ (whether pollution also makes crashes more severe conditional on their occurring) and places the pollution effect in context relative to other environmental factors.

To examine the determinants of crash severity, I estimate an ordered logit model using individual-level crash data. The dependent variable is the injury severity category assigned to each crash, which takes five ordered values: No Injury, Reported Not Evident, Non-Incapacitating Injury, Incapacitating Injury, and Fatal. These categories are described in Table 2. The independent variables include the posted speed limit, lighting conditions (daylight, dawn, dusk, dark with lighting, dark without lighting), weather conditions at the time of the crash, road surface conditions, and an indicator for whether the crash occurred during rush hour (7–9am or 4–6pm). I also include month and year fixed effects to control for seasonal and temporal variation. All estimates are conditional on a crash having occurred.

The ordered logit model is the appropriate choice for this outcome because injury severity has a natural ordering (No Injury, Reported Not Evident, Non-Incapacitating, Incapacitating, and Fatal) that OLS would ignore. OLS treats the difference between adjacent categories as equal in magnitude, implying that the gap between No Injury and Reported Not Evident is the same as the gap between Incapacitating and Fatal. Ordered logit respects the ranking without imposing equal spacing between categories.

Model fit is assessed using the Akaike Information Criterion (AIC), where lower values indicate better fit. Adding weather, lighting, and road surface controls in Column (2) reduces the AIC from 1,013,432 to 1,004,019 (an improvement of over 9,400 points), indicating that these observable road and environmental conditions explain a meaningful share of the variation in crash severity. Adding the rush hour indicator in Column (3) yields a further modest improvement of 28 points. The remaining unexplained variation reflects that many factors are driving crash severity.

The ordered logit model is estimated as follows. Let $Y_i \in \{1, 2, 3, 4, 5\}$ denote the injury severity category for crash i , where 1 = No Injury, 2 = Reported Not Evident, 3 = Non-Incapacitating, 4 = Incapacitating, and 5 = Fatal. The model assumes that there exists a latent continuous variable Y_i^* representing the underlying severity of the crash:

$$Y_i^* = X_i' \beta + \epsilon_i$$

where X_i is a vector of crash-level covariates and ϵ_i follows a logistic distribution. The observed severity category is determined by unknown threshold parameters $\tau_1 < \tau_2 < \tau_3 < \tau_4$:

$$Y_i = k \quad \text{if} \quad \tau_{k-1} < Y_i^* \leq \tau_k, \quad k = 1, \dots, 5$$

The cumulative probability that crash i falls in category k or below is:

$$\Pr(Y_i \leq k \mid X_i) = \frac{1}{1 + \exp(-(\tau_k - X_i' \beta))}$$

The model is estimated by maximum likelihood. The coefficients β are reported as log-odds ratios: a positive coefficient means that a unit increase in the corresponding variable increases the log-odds of being in a higher severity category. Exponentiating the coefficient gives the odds ratio: $\exp(\hat{\beta}_j)$ is the multiplicative change in the odds of being in a higher severity category for a one-unit increase in X_j . For example, the coefficient on Posted Speed Limit in Column (1) is 0.046, implying $\exp(0.046) \approx 1.047$, or a 4.7% increase in the odds of more severe injury for each additional mph of posted speed limit. Throughout the discussion, coefficients are converted to percentage changes in odds using this formula for ease of interpretation. The threshold parameters τ_k are estimated jointly with β but are not reported in the tables. These are descriptive regressions of crash severity on observable road, weather, lighting, and road surface conditions rather than a causal analysis.

The results explain several patterns in the descriptive statistics (Table 4). Posted speed limits have a positive and significant relationship with injury severity across all specifications. In the baseline model (Column 1), each additional mile per hour of posted speed limit increases the log odds of more severe injury by 0.046. This coefficient remains stable when

weather, lighting, and road surface controls are added (0.043 in Columns 2 and 3), with t-statistics exceeding 59 in all specifications. It indicates that each 1 mph increase in posted speed raises the odds of being in a higher severity category by 4.4%.

Lighting conditions also have effects on crash severity. Relative to daylight crashes, those that occur on dark but lighted roads show substantially higher severity, with a coefficient of 0.313, meaning crashes on dark, lighted roads have 37% higher odds of more severe injury compared to daylight crashes. Dawn and dusk periods also show increased severity, with coefficients of 0.180 and 0.121 respectively. Interestingly, crashes on unlighted dark roads show no significant difference from daylight crashes. This counterintuitive result may suggest that unlighted roads tend to be lower-speed residential streets where drivers exercise greater caution in darkness.

Weather conditions show a less straightforward pattern. Severe weather conditions exhibit the strongest effect on crash severity, corresponding to a 25% increase in the odds of more severe injury. Freezing rain also shows higher severity, corresponding to an 18% increase. These patterns are consistent with drivers potentially reducing speed and being cautious during obviously hazardous conditions like snow, but failing to adjust sufficiently for more subtle hazards like wet pavement or reduced visibility in rain.

Road surface conditions reinforce this interpretation. Wet surfaces increase crash severity, while ice or snow on the roadway shows a large negative coefficient, implying a 28% reduction in the odds of severe injury. This suggests that drivers respond more cautiously to visible ice and snow than to wet pavement, which may be less visibly apparent but still worsens vehicle handling.

The rush hour coefficient in Column (3) is negative and statistically significant. Crashes occurring during peak commute times (7–9am or 4–6pm) have 3.4% lower probability of severe injury compared to off-peak crashes. This finding aligns with the descriptive evidence showing that fatal crashes are disproportionately concentrated outside rush hours. Traffic congestion during peak periods likely constrains vehicle speed, reducing crash severity even as crash frequency may be higher due to increased traffic volume.

The ordered logit estimates show that crash severity is related to observable factors of the road environment. Higher speed, reduced lighting, certain weather conditions, and off-peak timing all contribute to more severe injuries conditional on a crash occurring. These findings provide a foundation for examining whether air pollution, through its documented effects on cognitive function, affects crash severity.

I now examine whether air pollution affects injury outcomes conditional on a crash occurring (Table 5).¹ If PM_{2.5} impairs cognitive function, it might both increase crash likelihood and also affect drivers' ability to avoid or mitigate collisions, potentially leading to more severe injuries.

Table 5 presents ordered logit estimates examining the relationship between PM_{2.5} concentrations and crash severity. Column (1) reproduces the baseline specification from the previous section for comparison, excluding pollution measures. Column (2) adds daily average PM_{2.5} concentration. Column (3) adds year and month fixed effects to control for seasonal patterns.

¹An IV version of the ordered logit is not reported because the ordered logit estimates are intentionally conditional on a crash having occurred and are meant to be descriptive of the severity rather than causal claims about PM_{2.5} exposure. The causal evidence for pollution's effect on both crash frequency and the share of fatal outcomes is provided by the aggregate IV results in Section 3.

The results show a small but statistically significant association between air pollution and crash severity when temporal fixed effects are included. In the specification without fixed effects (Column 2), the PM2.5 coefficient is small and statistically insignificant, but once year and month fixed effects are added in Column (3), the coefficient becomes significant. This suggests that the relationship between pollution and severity is confounded by seasonal factors. Both PM2.5 concentrations and crash severity show seasonal variation: pollution tends to be higher in winter and summer months due to heating and traffic patterns, while crash severity varies with seasonal changes in driving conditions and traffic.

The estimated coefficient in Column (3) implies that a 1 $\mu\text{g}/\text{m}^3$ increase in PM2.5 raises the odds of being in a more severe injury category by 0.3%. For a one standard deviation increase in PM2.5 (approximately 4.7 $\mu\text{g}/\text{m}^3$), the odds of more severe injury increase by about 1.4%. This effect is smaller than the effects of other environmental factors. For comparison, crashes occurring on dark but lighted roads have 39% higher odds of severe injury relative to daylight crashes, and crashes on wet surfaces have 19% higher odds compared to dry conditions.

A key finding of this analysis is that air pollution has a larger percentage effect on fatal crashes (9% per $\mu\text{g}/\text{m}^3$ in the 4-bin specification) compared to total crashes (0.78% per $\mu\text{g}/\text{m}^3$). This differential effect suggests that PM2.5 exposure affects not only the likelihood of a crash occurring but also the severity of crashes that do occur. This finding connects directly to the ordered logit results in Table 5, which show that PM2.5 increases crash severity conditional on a crash occurring. The ordered logit estimates indicate that a 1 $\mu\text{g}/\text{m}^3$ increase in PM2.5 increases the odds of being in a more severe injury category by

0.3%. When combined with an increase in the overall number of crashes, the result is a disproportionately large increase in fatal outcomes.

The mechanism linking pollution to crash severity likely operates through cognitive impairment. Previous research has established that PM_{2.5} exposure impairs cognitive function, including attention, reaction time, and decision-making. The crash severity analysis (Tables 4–5) showed that factors requiring quick reactions, such as driving in darkness or at higher speeds, are strongly associated with more severe outcomes. If pollution impairs drivers’ cognitive function, affected drivers may be less able to avoid crashes entirely (increasing crash frequency) and less able to take evasive action or mitigate impact when crashes occur (increasing crash severity).

Supporting this interpretation, the ordered logit results show that pollution’s effect on severity only becomes significant when year and month fixed effects are included (Table 5, Column 3), highlighting the importance of controlling for seasonal confounders in both the individual-level severity analysis and the aggregate crash count analysis.

The ordered logit estimates in Table 5 should be interpreted as descriptive. Like OLS, the ordered logit is subject to both omitted variable bias and measurement error. The measurement error in PM_{2.5} is likely non-classical. Monitors are placed at fixed locations and may systematically over- or understate exposure for drivers located far from the nearest monitor. It would mean the gap between measured and true exposure is correlated with location and time of day rather than random. The measurement error does not necessarily attenuate estimates toward zero, making the direction of bias ambiguous. On the omitted variable side, the bias could run in either direction too. High pollution days tend to coincide

with heavy traffic and economic activity, which independently increase crash severity, pushing the estimate upward. At the same time, conditions that generate high pollution may reduce driving speeds and alter the mix of drivers on the road, pushing the estimate downward. The inclusion of month-by-year fixed effects in Column (3) absorbs some of this seasonal confounding, which is consistent with the coefficient becoming significant only after their inclusion.

Several factors may explain why the Chicago estimates are substantial. First, Chicago is a dense urban environment with high traffic volume, where cognitive impairment may have larger consequences due to more complex driving conditions and greater interaction between vehicles. Second, the estimates here focus on a single city with relatively uniform monitoring, potentially reducing measurement error compared to county-level averages across diverse geographic areas.

2.5 Conclusion

This paper provides causal evidence that short-term exposure to fine particulate matter ($\text{PM}_{2.5}$) significantly increases traffic crashes in Chicago. Using an instrumental variables approach that exploits exogenous variation in wind direction, I find that a $1 \mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ leads to a 0.78% increase in total crashes and a 9.05% increase in fatal crashes. Pollution affects both the frequency of crashes and their severity when they occur.

The estimated effects are broadly consistent with related work. The first chapter of this dissertation finds a 1.34% increase in fatal crashes nationally per $1 \mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$,

while [Sager \(2019\)](#) finds a 2% increase in total accidents in the UK. The present paper finds a 0.78% increase in total crashes and a 9.05% increase in fatal crashes for Chicago, with the larger effect on fatal crashes reflecting pollution’s additional impact on crash severity conditional on a crash occurring.

The analysis reveals that pollution’s impact operates through multiple channels. At the individual level, $\text{PM}_{2.5}$ increases crash severity conditional on a crash occurring, with a 1 $\mu\text{g}/\text{m}^3$ increase in pollution raising the odds of being in a more severe injury category by 0.3%. This effect is smaller than the impacts of factors such as darkness (37% higher odds) or wet road surfaces (19% higher odds), but it provides direct evidence that pollution impairs the cognitive functions necessary for crash avoidance. At the aggregate level, pollution increases the total number of crashes, suggesting impaired attention leads to more frequent driving errors. Together, these effects produce a large increase in fatal outcomes: when crashes occur under higher pollution conditions, drivers are both more likely to crash and less able to take evasive action.

These findings have important implications for policy evaluation and public health. First, they highlight a significant external cost of air pollution that is typically omitted from regulatory cost-benefit analyses. Standard evaluations of air quality regulations focus on direct health outcomes such as cardiovascular and respiratory mortality. However, the traffic safety benefits of pollution reduction are substantial. For Chicago specifically, a one standard deviation reduction in $\text{PM}_{2.5}$ (4.66 $\mu\text{g}/\text{m}^3$) would prevent approximately 47 traffic fatalities per year. Using the Environmental Protection Agency’s value of a statistical life of \$7.4 million, this translates to roughly \$348 million in annual benefits from reduced traffic deaths alone.

These benefits are in addition to the well-documented health improvements from reduced respiratory and cardiovascular disease. The full social value of air quality improvements may be considerably larger than current estimates suggest.

To make these estimates more interpretable, consider the distribution of pollution days in Chicago. Over the sample period, approximately 500 days (roughly 63 days per year) exceeded the EPA’s “Moderate” threshold of $12 \mu\text{g}/\text{m}^3$. On a day where $\text{PM}_{2.5}$ reaches $15 \mu\text{g}/\text{m}^3$ ($3 \mu\text{g}/\text{m}^3$ above the Moderate threshold), the estimates imply approximately 7 additional total crashes and 0.08 additional fatal crashes relative to a clean air day. Policies that bring Chicago into compliance with the EPA annual standard of $12 \mu\text{g}/\text{m}^3$ on all days would therefore generate substantial traffic safety benefits in addition to the well-documented health gains.

Second, the results suggest potential interventions that could save lives before air quality improvements are achieved. Cities already issue air quality alerts when pollution reaches unhealthy levels. These alerts typically focus on health risks for vulnerable populations such as elderly individuals and those with respiratory conditions. The findings here suggest that high pollution days also pose higher crash risks for all drivers. Traffic safety campaigns could incorporate air quality warnings, alerting drivers that cognitive performance may be impaired on high pollution days and encouraging extra caution. Transportation agencies could implement dynamic speed limits or issue recommendations. Unlike long-term pollution reduction strategies that require years of infrastructure investment and regulatory change, such interventions could be implemented immediately.

Third, the disproportionate effect on fatal crashes compared to total crashes suggests

that pollution’s cognitive impacts are consequential. The individual-level severity analysis reveals that factors requiring quick reactions (such as driving in darkness or at higher speeds) are strongly associated with more severe outcomes. If pollution impairs reaction times and decision-making capacity, drivers may be unable to respond effectively when unexpected hazards arise. This interpretation is consistent with the literature documenting pollution’s effects on cognitive function.

The results raise several directions for future research. First, do pollution’s effects vary by driver characteristics? Professional drivers such as taxi and rideshare operators experience greater exposure and may exhibit different responses than occasional drivers. Older adults and individuals with pre-existing health conditions may be more sensitive to pollution’s cognitive effects. Second, do pollution’s effects differ across crash types? Multi-vehicle crashes at intersections may be more sensitive to impaired attention, while single-vehicle crashes on highways may be more sensitive to impaired reaction times. Third, how do pollution’s effects vary across cities with different traffic patterns and demographic compositions? Chicago is a dense urban environment with high traffic volume; the effects may differ in less congested settings or cities with different pollution sources. Addressing these questions would deepen our understanding of the mechanisms linking air quality to traffic safety and help target interventions most effectively. Finally, future work could examine whether pollution’s effects on severity interact with other crash risk factors, such as whether cognitive impairment is more consequential at higher speeds or in low-visibility conditions.




Figure 2.1: Total Daily Crashes in Chicago




Figure 2.2: Average Crashes by Month




Figure 2.3: Crash counts by hour of the day in Chicago.




Figure 2.4: Crash counts by day of the week in Chicago.




Figure 2.5: Average PM2.5 by Month




Figure 2.6: Daily PM2.5 Concentration in Chicago




Figure 2.7: First Stage: Mean PM2.5 Concentration by Wind Direction

Table 2.1: Descriptive Statistics

Variable	N	Mean	Std	Min	25%	50%	75%	Max
<i>Traffic Crashes (daily)</i>								
Total Crashes	2,922	291.77	63.95	85.00	255.00	298.50	333.00	583.00
Fatal Crashes	2,922	0.33	0.60	0.00	0.00	0.00	1.00	6.00
Non-Fatal Crashes	2,922	290.79	63.82	84.00	254.00	297.00	331.00	582.00
<i>Weather Conditions (daily)</i>								
Avg. Temperature (°C)	2,576	12.63	10.16	−25.87	4.44	13.39	21.73	31.71
Avg. Rain Intensity (mm/hr)	2,576	0.15	0.49	0.00	0.00	0.00	0.04	6.22
Avg. Total Rain (mm)	2,576	123.77	147.97	0.00	20.38	57.99	177.51	912.20
Avg. Wind Direction (deg.)	2,576	177.26	70.13	6.00	121.79	177.06	233.46	338.46
Avg. Wind Speed (m/s)	2,576	2.72	1.62	0.58	1.79	2.41	3.28	44.83
Max. Wind Speed (m/s)	2,576	10.33	35.03	1.70	6.50	8.50	11.00	—
<i>Air Quality (daily)</i>								
Avg. PM _{2.5} (µg/m ³)	2,650	9.18	4.66	0.40	6.10	8.50	11.30	81.00
Avg. AQI	2,650	44.95	14.85	2.00	34.00	47.00	55.00	169.00

Notes: Descriptive statistics for daily observations of traffic crashes, weather conditions, and air pollution levels in Chicago, 2017–2025.

Table 2.2: Crash Characteristics by Injury Severity

	Injury Severity Category				
	No Injury	Reported Not Evident	Non-Incapacitating	Incapacitating	Fatal
Observations					
Count	816,068	42,652	75,364	15,680	1,024
Percent of total	85.83	4.49	7.93	1.65	0.11
Roadway Characteristics					
Avg. speed limit (mph)	28.2	29.6	29.6	29.8	30.0
Temporal Patterns (%)					
Rush hour (7–9am, 4–6pm)	35.7	36.3	33.6	30.7	16.2
Environmental Conditions (%)					
Darkness (lighted or unlighted)	25.6	28.1	33.1	36.4	57.7
Adverse weather	18.9	18.0	16.2	15.1	15.2
Wet or icy road surface	16.5	19.9	18.8	17.7	17.2

Notes: Summary statistics for all reported crashes in Chicago, 2017–2025, by officer-assigned injury severity. Percentages are shares within each severity column.

Table 2.3: Distribution of Weather and Lighting Conditions

Condition	Observations	Percent (%)
Weather Condition		
Clear	747,397	78.4
Rain	82,247	8.6
Snow	30,715	3.2
Cloudy/Overcast	28,128	3.0
Unknown	58,880	6.2
Freezing Rain/Sleet	3,435	0.4
Fog/Smoke/Haze	1,413	0.1
Severe Weather	671	0.1
Lighting Condition		
Daylight	611,168	64.1
Darkness, lighted road	208,329	21.9
Unknown	45,786	4.8
Darkness, unlighted	44,710	4.7
Dusk	27,056	2.8
Dawn	15,837	1.7

Notes: Distribution of weather and lighting conditions across all crashes in the sample (952,886 total crashes). The majority occur in clear weather (78.4%) and daylight (64.1%). Severe weather includes blowing snow, blowing sand/soil/dirt, and severe wind.

Table 2.4: Ordered Logit: Determinants of Crash Injury Severity

	Dependent variable: Injury Severity		
	(1)	(2)	(3)
Roadway Characteristics			
Posted Speed Limit	0.046*** (0.001)	0.043*** (0.001)	0.043*** (0.001)
Weather Conditions			
Cloudy		0.113*** (0.017)	0.113*** (0.017)
Rain		0.029 (0.016)	0.030 (0.016)
Snow		-0.061** (0.024)	-0.061** (0.024)
Freezing Rain		0.164*** (0.048)	0.164*** (0.048)
Fog		0.104 (0.070)	0.105 (0.070)
Severe Weather		0.225** (0.108)	0.226** (0.108)
Unknown		-0.251*** (0.024)	-0.251*** (0.024)
Lighting Conditions			
Dawn		0.183*** (0.022)	0.180*** (0.022)
Dusk		0.117*** (0.017)	0.121*** (0.018)
Darkness, Lighted Road		0.322*** (0.007)	0.313*** (0.007)
Darkness, Unlighted		0.016 (0.015)	0.009 (0.015)
Unknown		-0.809*** (0.028)	-0.811*** (0.028)
Road Surface			
Wet		0.121*** (0.014)	0.122*** (0.014)
Ice/Snow		-0.325*** (0.023)	-0.323*** (0.023)
Other		0.225*** (0.054)	0.225*** (0.054)
Unknown		-0.421*** (0.017)	-0.421*** (0.017)
Temporal Controls			
Rush Hour			-0.035*** (0.006)
Observations	915,224	915,224	915,224
AIC	1,013,432	1,004,019	1,003,991

Notes: *p<0.10; **p<0.05; ***p<0.01. Standard errors in parentheses. Dependent variable is injury severity (ordered): None, Reported Not Evident, Non-Incapacitating, Incapacitating, Fatal. Coefficients are log odds ratios from ordered logistic regression. Reference: clear weather, daylight, dry road. Column (1) includes only posted speed limit; (2) adds weather, lighting, and surface conditions; (3) adds rush hour indicator. Sample excludes crashes with missing severity or speed limits below 15 or above 70 mph.

Table 2.5: Ordered Logit: Effect of PM2.5 on Crash Injury Severity

	Dependent variable: Injury Severity		
	Base (1)	Add PM2.5 (2)	Add FE (3)
Air Quality			
PM2.5 ($\mu\text{g}/\text{m}^3$)		0.0005 (0.001)	0.003*** (0.001)
Roadway Characteristics			
Posted Speed Limit	0.044*** (0.001)	0.044*** (0.001)	0.044*** (0.001)
Weather Conditions			
Cloudy	0.096*** (0.019)	0.096*** (0.019)	0.111*** (0.019)
Rain	0.033 (0.017)	0.033 (0.017)	-0.007 (0.018)
Snow	-0.041 (0.026)	-0.040 (0.026)	0.002 (0.027)
Freezing Rain	0.171*** (0.054)	0.171*** (0.054)	0.160*** (0.054)
Lighting Conditions			
Dawn	0.178*** (0.024)	0.178*** (0.024)	0.190*** (0.024)
Dusk	0.129*** (0.019)	0.129*** (0.019)	0.146*** (0.019)
Darkness, Lighted Road	0.312*** (0.008)	0.311*** (0.008)	0.332*** (0.008)
Darkness, Unlighted	0.014 (0.016)	0.014 (0.016)	0.048*** (0.016)
Road Surface			
Wet	0.111*** (0.015)	0.111*** (0.015)	0.170*** (0.015)
Ice/Snow	-0.337*** (0.026)	-0.337*** (0.026)	-0.197*** (0.026)
Temporal Controls			
Rush Hour	-0.033*** (0.007)	-0.033*** (0.007)	-0.021*** (0.007)
Year FE	No	No	Yes
Month FE	No	No	Yes
Observations	747,431	747,431	747,431
AIC	838,570	838,572	836,318

Notes: * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$. Standard errors in parentheses. Dependent variable is injury severity (ordered): None, Reported Not Evident, Non-Incapacitating, Incapacitating, Fatal. Coefficients are log odds ratios from ordered logistic regression. Reference: clear weather, daylight, dry road. Column (1) excludes PM2.5; (2) adds PM2.5; (3) adds year and month FE. Sample: 2017-2024 (excluding 2020-2021), complete data, speed limits 15-70 mph. This examines reduced-form relationship; we do not instrument in individual-level setting due to complexity of IV with ordered models and because main causal identification comes from IV on aggregate crash counts.

Table 2.6: OLS: Effect of PM_{2.5} on Daily Crashes in Chicago

	Dependent Variable		
	Total Crashes (1)	Fatal Crashes (2)	Non-Fatal Crashes (3)
PM _{2.5} ($\mu\text{g}/\text{m}^3$)	−0.080 (0.263)	−0.001 (0.003)	−0.079 (0.262)
Rain Intensity	10.055*** (2.525)	0.010 (0.027)	9.866*** (2.519)
Wind Speed	−2.163*** (0.741)	−0.020** (0.008)	−2.162*** (0.739)
Day-of-Week FE	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes
Weather Controls	Yes	Yes	Yes
Mean of Dep. Var.	297.0	0.305	296.0
Observations	1,924	1,924	1,924

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors in parentheses. Weather controls include rain intensity and wind speed.

Table 2.7: IV: Effect of PM_{2.5} on Daily Crashes in Chicago (4 Wind Direction Bins)

	Dependent Variable		
	Total Crashes (1)	Fatal Crashes (2)	Non-Fatal Crashes (3)
PM _{2.5} ($\mu\text{g}/\text{m}^3$)	2.314** (1.134)	0.028** (0.012)	2.300** (1.131)
Rain Intensity	10.600*** (2.591)	0.016 (0.028)	10.407*** (2.585)
Wind Speed	-1.061 (0.911)	-0.007 (0.010)	-1.066 (0.909)
Day-of-Week FE	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes
Weather Controls	Yes	Yes	Yes
Number of Instruments	3	3	3
Mean of Dep. Var.	297.0	0.305	296.0
Percentage Effect	0.78%	9.05%	0.78%
First-stage F-stat	37.8	37.8	37.8
Observations	1,924	1,924	1,924

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors in parentheses. Instruments: Three wind direction bin indicators (90° bins: N-NE, E-SE, S-SW; reference: W-NW). Weather controls include rain intensity and wind speed. First-stage F-statistic exceeds Stock-Yogo critical value of 22.3 for 10% maximal IV size with three instruments.

Table 2.8: IV: Effect of PM_{2.5} on Daily Crashes in Chicago (8 Wind Direction Bins)

	Dependent Variable		
	Total Crashes (1)	Fatal Crashes (2)	Non-Fatal Crashes (3)
PM _{2.5} ($\mu\text{g}/\text{m}^3$)	1.600* (1.079)	0.020* (0.012)	1.593* (1.076)
Rain Intensity	10.437*** (2.563)	0.015 (0.028)	10.246*** (2.557)
Wind Speed	-1.389 (0.890)	-0.010 (0.010)	-1.392 (0.888)
Day-of-Week FE	Yes	Yes	Yes
Month-Year FE	Yes	Yes	Yes
Weather Controls	Yes	Yes	Yes
Number of Instruments	7	7	7
Mean of Dep. Var.	297.0	0.305	296.0
Percentage Effect	0.54%	6.46%	0.54%
First-stage F-stat	17.6	17.6	17.6
Observations	1,924	1,924	1,924

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors in parentheses. Instruments: Seven wind direction bin indicators (45° bins; reference: NW 315°–360°). Weather controls include rain intensity and wind speed. First-stage F-statistic is below Stock-Yogo critical value of 22.3, suggesting potential weak instrument concerns. Results presented for robustness; main specifications use 4-bin instrument.

Chapter 3:

Effect of the Clean Air Act (1970) on Mental Health Outcomes

Irina Firsova

University of California, Davis

Abstract

This paper estimates the effect of the 1970 Clean Air Act Amendments on mental health outcomes, using county-level suicide rates as the primary outcome measure. Exploiting variation in total suspended particulates (TSP) nonattainment status as an instrumental variable for changes in air pollution between 1970 and 1980, I estimate first-differenced equations that absorb time-invariant county characteristics. Two-stage least squares estimates are small in magnitude and statistically insignificant across all specifications, with point estimates implying that the average TSP reduction generated by Clean Air Act enforcement would be associated with a change in the age-adjusted suicide rate of at most 0.50 per 100,000. Population-weighted estimates, which give greater influence to large urban counties, yield a positive and significant effect of TSP on suicide rates, suggesting that the effect may be concentrated in large urban counties. The null result is robust to alternative instrument definitions and to restricting the sample to counties near the nonattainment threshold. Heterogeneity analysis by race and gender yields uniformly insignificant estimates across all demographic subgroups, with the largest point estimates found for white males.

3.1 Introduction

Depression and suicide impose an enormous burden on individuals and societies worldwide, yet the environmental determinants of mental health remain poorly understood. A growing body of evidence suggests that air pollution may be one such determinant, operating through both biological pathways (including neuroinflammation and disruption of neurotransmitter systems) and socioeconomic channels such as reduced labor productivity and income loss. In the United States alone, the economic burden of depression, including direct medical costs and workplace costs, was estimated at \$210.5 billion in 2010 ([Greenberg, Fournier, Sisitsky, Pike, and Kessler, 2015](#)). At the global level, depression and anxiety cost the global economy an estimated \$1 trillion per year in lost productivity ([World Health Organization, 2025](#)). This paper asks whether reductions in air pollution causally improve mental health outcomes, exploiting the regulatory variation generated by the 1970 Clean Air Act Amendments to estimate the effect of particulate matter on county-level suicide rates.

Depression affects approximately 350 million people worldwide, representing 6% of the global population ([Vos, Allen, Arora, Barber, Bhutta, Brown, Carter, Casey, Charlson, Chen et al. \(2016\)](#)), and is associated with substantially lower quality of life, elevated cardiovascular mortality, and increased suicide risk ([Ruo, Rumsfeld, Hlatky, Liu, Browner, and Whooley \(2003\)](#)). The existing medical literature documents associations between air pollution and poor mental health outcomes, but these studies largely rely on observational designs that cannot establish causality. For example, [Szyszkowicz, Willey, Grafstein, Rowe, and Colman \(2010\)](#) find that suicide attempts increase during periods of elevated particulate matter

concentrations. Studies examining the timing of pollution exposure relative to suicide find that PM10 concentrations have the strongest association in the 0 to 2 days prior to the event, while PM2.5 concentrations are most strongly associated one day prior (Go, Kim, Choi, Han, and Kim (2024)). Evidence from Guangzhou, China indicates that particulate matter is associated with higher suicide rates, with effects that vary across social groups, are stronger during winter months, and are more pronounced among men with higher levels of education (Lin, Li, Song, Zhou, Shen, and Ou (2016)). Separately, Lundberg, Cantor-Graae, Rukundo, Ashaba, and Östergren (2009) finds that mental health outcomes are worse in urban areas than in rural ones, primarily to urban-associated stress, though lower air quality in cities may also contribute. The inability of these studies to isolate causal effects from confounding factors, such as socioeconomic conditions that jointly determine both pollution exposure and mental health, motivates the quasi-experimental approach employed in this paper.

From a policy perspective, accurate estimation of the health costs of air pollution is essential for designing socially efficient environmental regulations, where both costs and benefits are fully considered. If pollution imposes substantial mental health care costs, and these are ignored, then the social cost of pollution will be underestimated, and optimal pollution taxes will be set too low. This concern is driven by evidence that the negative physical health effects of air pollution are heterogeneous across income groups (Jans, Johansson, and Nilsson (2018)): lower-income populations bear a disproportionate share of the health burden, making mental health impacts of pollution possibly an important aspect of environmental health inequality.

The effects of air pollution on mental health remain understudied, despite reasonable theoretical grounds for expecting such a relationship. Depression is associated with social class, income, health status, and living conditions ([Lehtinen and Joukamaa \(1994\)](#)), and loss of income and unemployment are well-established risk factors. Reductions in labor productivity and financial burdens arising from pollution-induced illness are therefore plausible pathways through which air pollution could contribute to depressive episodes.

A substantial body of literature documents that air pollution leads to negative health outcomes more broadly, primarily through respiratory and cardiovascular disease ([Kim et al. \(2013\)](#), [Fiordelisi et al. \(2017\)](#)). The resulting chronic illness and disability impose a considerable burden on public health expenditures ([Segalowitz \(2008\)](#)). The negative effects of air pollution on cognition are also well-documented. Numerous studies find that exposure to air pollutants reduces labor productivity through cognitive decline ([Hausman et al. \(1984\)](#), [Carson et al. \(2011\)](#), [Rodrigues-Silva et al. \(2012\)](#), [Hanna and Oliva \(2015\)](#)).

This paper contributes to a growing body of quasi-experimental work estimating the causal effects of air pollution on health and economic outcomes. Early contributions in this literature, including [Chay and Greenstone \(2005\)](#), established that regulatory variation generated by the Clean Air Act could be exploited for causal identification, and demonstrated that TSP reductions led to meaningful improvements in infant mortality and housing values. Subsequent work has extended this framework to a range of outcomes, documenting effects of pollution exposure on cognitive performance, educational attainment, labor productivity, cardiovascular health, and long-run labor market outcomes ([Isen, Rossin-Slater, and Walker \(2017\)](#)). [Currie, Voorheis, and Walker \(2020\)](#) exploit satellite-based measures of air quality

to document racial disparities in pollution exposure and trace their origins to Clean Air Act enforcement. They highlight the distributional dimension of environmental regulation that motivates the heterogeneity analysis in this paper. Relative to this literature, this paper offers a county-level analysis of mental health outcomes over the first decade of Clean Air Act enforcement.

The 1970 Clean Air Act Amendments provide a good setting for causal identification of the effects of air pollution on health outcomes. Prior to the Act, air quality regulation was largely left to individual states, and very few chose to impose meaningful restrictions. The federal legislation introduced uniform standards and established a county-level compliance system administered by the Environmental Protection Agency, creating discontinuities in regulatory intensity across otherwise similar counties. Counties exceeding TSP concentration thresholds were assigned nonattainment status and subjected to substantially more stringent emissions requirements than attainment counties, generating large, policy-driven reductions in particulate matter that were plausibly independent of local economic conditions. Because nonattainment designation was determined by whether measured TSP concentrations crossed a federally mandated threshold, rather than by local preferences or economic circumstances, it provides a source of exogenous variation in air quality that can be used to isolate the causal effect of pollution on suicide rates from confounding factors that jointly determine both pollution levels and mental health outcomes.

This paper uses county-level suicide mortality as its measure of mental health and finds null effects of Clean Air Act-induced TSP reductions on suicide rates. However, population-weighted estimates suggest that the effect may be present in large urban counties, where

pollution exposure is higher and more consistently measured. Suicide rates serve as the primary outcome measure in this paper for several reasons. Mental health outcomes are notoriously difficult to measure at the population level, as the most commonly used instruments rely on self-reported symptoms and are rarely available for large, representative samples over long time horizons. Administrative mortality data, by contrast, provides comprehensive county-level coverage across the full study period, making it well-suited to the research design employed here. Suicide is also a severe and consequential endpoint that represents mental health deterioration. There is established evidence linking depression to elevated suicide risk ([Ruo et al. \(2003\)](#)). That said, suicide mortality is an imperfect proxy for mental health more broadly, as it reflects only the most extreme outcomes and may therefore understate the true effect of air pollution on the distribution of mental health in the population. Misclassification is also a concern: the social stigma associated with suicide creates incentives for underreporting, and classification practices varied across jurisdictions and over time during the study period ([Rockett, Kapusta, Bhandari et al. \(2011\)](#)). These limitations should be kept in mind when interpreting the null results reported below, as they imply that the absence of a detectable effect on suicide rates does not rule out subtler effects on mental health more broadly.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of the Clean Air Act and the regulatory framework governing TSP nonattainment status. Section 3 describes the data and empirical framework. Section 4 presents the identification strategy and empirical results. Section 5 examines heterogeneity by race and gender. Section 6 discusses limitations, and Section 7 concludes the paper.

3.2 Institutional background

3.2.1 Background on the Clean Air Act

Growing concern in Congress throughout the 1960s over the health impacts of worsening urban air quality led to the 1970 Clean Air Act Amendments. It was the first attempt to federally regulate air pollution in a comprehensive way. Prior to its introduction, it was up to states to decide whether to impose such regulations, and very few states actually did. The Clean Air Act introduced the National Ambient Air Quality Standards for six criteria pollutants (particulate matter, ozone, sulfur dioxide, nitrogen dioxide, carbon monoxide, and lead) and used them to determine counties' compliance with the act. The Environmental Protection Agency (EPA) conducted an annual evaluation at the county level, and assigned nonattainment, or attainment status, to each county. Nonattainment status was assigned if either of the thresholds below was exceeded

- the annual geometric mean concentration exceeds $75\mu g/m^3$
- the second-highest daily concentration exceeds $260\mu g/m^3$

Under the Clean Air Act, all states were required to come up with plant-specific regulations for every major pollution source, with a focus on counties assigned nonattainment status, while attainment counties remained relatively unregulated. Effective regulations would require massive investments by already existing and new plants into cleaner technology, as well as the installation of pollution abatement equipment that monitors the emission levels

and sets a ceiling on them. Emission limits were determined by states.

Up until 1977 legislation expanding the Clean Air Act, the EPA had control over state plans in order to lower variability in requirements across the country, but it was states' responsibility to penalize counties for nonattainment. With the 1977 amendment, plant-specific regulations became the subject of federal law too, which gave the EPA more legal enforcement power.

Counties designated as nonattainment were required to submit State Implementation Plans (SIPs) to the EPA, detailing the specific measures by which they would bring pollution levels into compliance. For existing major stationary sources in nonattainment areas, SIPs were required to mandate the application of Reasonably Available Control Technology (RACT). This standard required meaningful and feasible reductions in emissions. New stationary sources in nonattainment areas were subject to New Source Performance Standards (NSPS). It imposed uniform emission limits. States that failed to submit approvable SIPs risked EPA intervention. It could lead to potential withholding of federal funds, giving states a strong fiscal incentive to comply. The practical effect was that nonattainment counties faced substantially higher regulatory burdens on industrial activity than attainment counties, generating the differential reductions in TSP concentrations that are the basis of the identification strategy in this paper.

3.2.2 Background on PM (particulate matter)

Particulate matter is made up of microscopic particles of solid and liquid matter suspended in the air. There are two types of particles that are regulated by the Environmental Protection Agency (EPA).

- PM10 particles are inhalable particles less than or equal to 10 micrometers in diameter. For example, dust, pollen, and mold are PM10 particles. These particles are small enough to get into the lungs. Exposures to PM10 have been linked to the worsening of respiratory diseases, such as asthma, leading to hospitalization and emergency department visits.
- PM2.5 particles are inhalable particles less than or equal to 2.5 micrometers in diameter. For example, combustion particles, metal, and organic compounds are PM2.5 particles. These particles are small enough to get into the lungs and enter the bloodstream. The main natural sources of PM2.5 are volcanic eruptions and wildfires. As for anthropogenic sources, they include fossil fuel combustion from power plants, industries, and automobiles. They can stay in the air for a long time and travel for hundreds of miles. Because they can penetrate buildings, many people are easily exposed to them. Long-term exposure to PM2.5 has been associated with premature death, particularly in people who have chronic heart or lung diseases, and reduced lung function growth in children.

Particulate matter of diameter between 2.5 and 10 micrometers is called “coarse particulates,”

while PM with diameter less than or equal to $2.5\text{ }\mu\text{m}$ is called “fine particulates.”

The biological pathways through which air pollution may affect mental health operate primarily through neuroinflammation and oxidative stress. Fine particulate matter, particularly PM_{2.5}, is small enough to enter the bloodstream and cross the blood-brain barrier, where it can trigger inflammatory responses in neural tissue ([Block and Calderon-Garciduenas \(2009\)](#)). Sustained neuroinflammation has been linked to dysregulation of neurotransmitter systems, including serotonin and dopamine, which are central to mood regulation and are determinants of depression, also contributing to chronic physiological stress that over time increases vulnerability to psychiatric illness ([Troubat, Barone, Leman et al. \(2021\)](#)). These mechanisms suggest that the relationship between air pollution and mental health may reflect a direct physiological pathway, independent of indirect socioeconomic channels.

3.3 Data

3.3.1 Air Pollution Data

The air pollution dataset I use was constructed by Chay and Greenstone (2005) from a combination of the 1972 and 1983 City and County Data Books, the EPA’s Air Quality Subsystem data file, and the Code of Federal Regulations. The data are measured at the county level and comprise 1,000 observations, corresponding to counties with particulate pollution monitors in place both at the beginning and end of the 1970s. These counties account for over 80% of the U.S. population, providing broad national coverage despite the

sample restriction.

The pollutant of interest is total suspended particulates (TSP). TSP refers to the total mass concentration of all airborne particles suspended in the air, regardless of size, and encompasses both PM10 and PM2.5. TSP is the most visible form of air pollution and, among the pollutants regulated under the Clean Air Act, has the most extensively studied health effects. The EPA monitoring infrastructure was also more developed for TSP than for other pollutants, as it is comparatively easier to measure, which contributes to the larger number of counties with pollution measurements available for this analysis.

County-level TSP concentrations are constructed from monitor-level data, consistent with the county-level unit at which EPA regulations are enforced. The measurements span the period 1969 to 1980. For counties with multiple monitors, the county mean is computed as a weighted average of monitor-specific geometric means (the geometric mean being the n -th root of the product of n observations, which reduces the influence of extreme values relative to the arithmetic mean), with the number of observations per monitor used as weights. To address measurement error, the 1970 TSP level for each county is defined as the average of non-missing annual geometric means over 1969–1972, and the 1980 level is similarly defined over 1977–1980. Averaging over multiple years also increases the number of counties with pollution measurements, as the number of active monitors increased over the decade. Since the empirical specification takes first differences between the 1970 and 1980 endpoints only, this multi-year averaging applies exclusively to the construction of those two endpoint measures and not to intermediate years.

A complication arises in the construction of the instrumental variable. The EPA did not

make nonattainment status assignments publicly available until 1978, so attainment status prior to that year must be approximated using underlying TSP concentration data and applicable state regulatory thresholds. This approximation introduces some imprecision into the instrument, though the strong first-stage relationship documented in Section 4 suggests this is not a significant source of bias.

A potential concern with the TSP data is that monitor placement was non-random. EPA criteria required monitors to capture the highest concentrations and representative population areas, which means monitored counties are not a random sample of all counties. Of the approximately 3,063 counties in the United States, only 1,000–1,400 had monitors, though these accounted for 80–85% of the US population. A related issue is monitor entry and exit. If new monitors are added in dirtier locations, measured county-level TSP changes may partially reflect compositional shifts in the monitor network rather than air quality changes. Following [Chay and Greenstone \(2005\)](#), the sample is restricted to counties with monitors present at both the beginning and end of the decade, which mitigates but does not fully resolve this concern. The regression discontinuity results in Section 4, which show jumps in pollution reduction at the regulatory threshold, provide evidence that monitor placement bias is not severe enough to confound the main estimates.

Counties without monitors cannot be assigned nonattainment status and are excluded from the analysis entirely. This means the sample is restricted to counties dirty enough to warrant monitoring in the first place. Those are systematically more industrial and more polluted than the average US county. As a result, attainment counties in the sample are not clean counties in general, but counties that fell below the regulatory threshold despite having

pollution levels high enough to justify a monitor. This could limit the external validity of the estimates. The results speak to the effect of TSP reductions in monitored counties, which account for 80–85% of the US population, but cannot be extrapolated to the full distribution of counties. Within this sample, the comparison between counties just above and just below the nonattainment threshold approximates a regression discontinuity design. It partially addresses the selection concern by focusing on counties that are otherwise similar in terms of pollution levels.

3.3.2 Suicide Rate Data

Suicide rate data are drawn from the CDC WONDER Compressed Mortality File (CMF), a county-level national mortality database covering the years 1968–2016, compiled from death certificates filed across all fifty states and the District of Columbia. The CMF provides counts of deaths, crude death rates, and age-adjusted death rates at the county level, with further disaggregation available by age, race, and gender. A key advantage of this data source is that it captures the full population rather than a sample, reducing concerns about sampling variability in county-level estimates.

Age-adjusted suicide rates are used as the primary outcome measure. Age adjustment is necessary because suicide risk varies across age groups, and counties differ in their age distributions. Without adjustment, differences in crude suicide rates across counties could reflect differences in age structure rather than true differences in suicide risk. Age adjustment is achieved by computing weighted averages of age-specific death rates using the US standard population as the fixed reference population. It allows for meaningful comparisons

of mortality risk across counties with different age distributions.

The cause of death recorded in the CMF is determined by the International Classification of Diseases (ICD) code listed in the underlying cause of death field on each death certificate. While this provides a standardized and consistently applied classification system, misreporting remains a concern. The social stigma associated with suicide creates incentives for families and medical examiners to record ambiguous deaths under alternative causes, and classification practices varied across jurisdictions during the study period. Misclassification and undercounting of suicide mortality are well-documented concerns in the epidemiological literature ([Rockett et al. \(2011\)](#)), and the extent of underreporting is likely to have been greater in the 1970s than in more recent decades. The implications of this for interpreting the results are discussed further in Section 6.

3.3.3 Descriptive Statistics

Table 1 reports descriptive statistics for the main variables. The mean change in TSP concentration between 1970 and 1980 is $-8.59 \mu\text{g}/\text{m}^3$, reflecting the substantial aggregate improvement in air quality following the introduction of the Clean Air Act. There is considerable heterogeneity across counties (standard deviation of $22.02 \mu\text{g}/\text{m}^3$). The distribution is left-skewed indicating that some counties experienced very large reductions, while the 75th percentile change of $3.81 \mu\text{g}/\text{m}^3$ suggests that a substantial share of counties saw little improvement. Approximately one third of counties in the sample were assigned nonattainment status, and as Table 2 documents, nonattainment counties experienced TSP reductions roughly $10 \mu\text{g}/\text{m}^3$ larger than attainment counties on average.

The mean change in the age-adjusted suicide rate over the same period is 0.46 per 100,000, with a standard deviation of 7.62, reflecting substantial dispersion across counties. The wide interquartile range, from -2.80 to 4.10 per 100,000, indicates that suicide rate trajectories varied considerably across the sample, driven by local economic and social conditions. Figure 1 illustrates the gender gap in age-adjusted suicide rates over the pre-reform period, with male rates consistently exceeding female rates throughout the 1968–1978 window. This disparity has implications for the heterogeneity analysis in Section 5, where the larger absolute levels of male suicide rates provide greater statistical power to detect effects for men than for women.

Table 3 reports the correlation matrix for the main variables. Changes in TSP are most strongly correlated with changes in manufacturing employment share and income per capita, consistent with the intuition that pollution levels are driven primarily by industrial activity and the economic conditions associated with it. The correlation between TSP changes and suicide rate changes is small and negative, which anticipates the counterintuitive OLS result discussed in Section 4 and motivates the instrumental variable approach. Changes in suicide rates show little correlation with the economic control variables, suggesting that the relationship between pollution and suicide rates is not straightforwardly driven by the observable economic shocks captured in the data.

3.3.4 Empirical Framework

This paper follows the identification strategy developed in [Chay and Greenstone \(2005\)](#), who exploit the regulatory structure of the Clean Air Act Amendments to estimate the causal

effect of air quality on housing values. Their approach takes advantage of the nonattainment designation to construct an instrumental variable for TSP changes, addressing the omitted variable bias that arises because pollution levels and economic conditions are jointly determined. They estimate first-difference equations for county-level housing prices between 1970 and 1980 and find that the elasticity of housing prices with respect to TSP levels lies within the interval $[-0.35, -0.2]$.

The present paper applies the same first-difference instrumental variable framework to estimate the effect of air pollution on suicide rates. The outcome here differs from that in [Chay and Greenstone \(2005\)](#): suicide rates are not a market measure, and there are no reliable estimates of the social costs of mental health that would allow for a welfare interpretation analogous to the hedonic price approach. The analysis therefore focuses on the absolute effect of TSP reductions on county-level suicide rates, rather than on the monetized welfare implications.

This paper applies the same first-difference instrumental variable framework to estimate the effect of air pollution on suicide rates. The design is appropriate here because first-differencing eliminates time-invariant county characteristics that together determine pollution and health outcomes. Also, nonattainment designation provides a source of exogenous variation in TSP changes that is orthogonal to local economic conditions. The outcome differs in that suicide rates are not a market measure, so the analysis focuses on the effect of TSP reductions on the change in county-level suicide rates rather than on monetized welfare implications.

3.3.5 Control Variables

The analysis incorporates a set of county-level control variables drawn from the 1972 and 1983 City and County Data Books including income per capita, unemployment, manufacturing employment share, population density, and urban share. Since all controls are measured as changes between 1970 and 1980, they capture within-county economic shifts over the decade, such as business cycle fluctuations and structural changes in the local economy, rather than permanent differences across places, which are absorbed by the first-difference design. The set of control variables captures the main economic and demographic channels through which omitted variable bias could affect the estimates. Changes in income per capita and unemployment rates are included because local labor market conditions are strongly associated with both industrial activity, and hence pollution levels, and mental health outcomes. Manufacturing employment share is included because counties with a larger industrial base tend to have higher TSP concentrations and may also have distinct social and occupational characteristics that affect suicide risk. Population density and urban share are included to absorb broad differences in the urban character of counties that correlate with both pollution exposure and access to mental health resources. As shown in Table 3, changes in TSP are most strongly correlated with changes in manufacturing employment and income, motivating their inclusion as controls in the main specifications.

3.4 Identification and Empirical Strategy

3.4.1 First-Difference Specification

The estimating equation takes the following first-difference form:

$$(Y_i^{80} - Y_i^{70}) = \mu + \beta(P_i^{80} - P_i^{70}) + \alpha(\mathbf{X}_i^{80} - \mathbf{X}_i^{70}) + (e_i^{80} - e_i^{70})$$

where Y_i is the county-level age-adjusted suicide rate, P_i is the geometric mean of TSP concentrations across all monitors in the county, $\mathbf{X}_i$ is a vector of county-level controls, and μ is a constant that captures the common nationwide time trend in suicide rates between 1970 and 1980. Taking first differences eliminates time-invariant county fixed effects, removing bias that would otherwise arise from permanent differences across counties in geography, institutions, or infrastructure that affect both pollution levels and mental health outcomes.

The parameter of interest is β , which captures the effect of a one unit change in mean TSP on the change in the suicide rate. The equation is estimated under three specifications: (1) no control variables, (2) adjusting for the main effects of the control variables, and (3) adjusting for main effects, polynomials, and interactions of all control variables in the dataset. The progression across specifications serves as a robustness check. It provides reassurance that the results are not sensitive to the particular set of controls included.

Identification of the OLS estimator of β requires the assumption that unobservable shocks

to suicide rates do not covary with unobservable shocks to air quality. This assumption is unlikely to hold in practice. Counties with higher TSP levels tend to be more urbanized and industrially active. These characteristics are associated with higher incomes but also with a range of social conditions that independently affect mental health. The correlation matrix in Table 3 confirms that TSP changes are correlated with changes in income and manufacturing employment, indicating that observable economic shocks move together with pollution levels. If unobservable shocks follow a similar pattern, OLS estimates of β will be biased.

3.4.2 OLS Results

Table 4 presents the OLS estimates. Across all three specifications, the estimated effect of TSP on suicide rates is negative and marginally significant at the 10% level. The point estimate implies that a 1 $\mu\text{g}/\text{m}^3$ increase in mean TSP is associated with a decrease of approximately 0.02 in the age-adjusted suicide rate. This result is counterintuitive: if air pollution worsens mental health, one would expect higher TSP levels to be associated with higher, not lower, suicide rates. The stability of the coefficient across specifications, and its small magnitude, suggest that the estimate is reflecting omitted variable bias rather than a causal relationship. As discussed above, areas with higher pollution tend to be more economically active, and the income and employment effects of industrial activity may be driving the negative association between TSP and suicide rates in the OLS estimates.

3.4.3 Instrumental Variable Strategy

To address the omitted variable bias in the OLS estimates, TSP changes are instrumented using a binary indicator for county-level nonattainment status. The instrument D_i takes the value 1 if county i was designated nonattainment in 1975–76 and 0 otherwise. Nonattainment status is time-varying: the EPA evaluates counties annually and counties can exit nonattainment once their TSP concentrations fall below the regulatory thresholds. In the pre-Act period, no federal standards existed and all counties were effectively in attainment. The nonattainment indicator therefore captures a change in regulatory status rather than a persistent pre-existing characteristic — since $D_i = 0$ for all counties before the Act, the instrument is analogous to a first difference, comparing counties that were pushed into nonattainment by the Act against those that were not.

The first-stage equation takes the form:

$$(P_i^{80} - P_i^{70}) = \mu + \gamma D_i + \delta(\mathbf{X}_i^{80} - \mathbf{X}_i^{70}) + \nu_i$$

where D_i is the nonattainment indicator, γ captures the effect of regulatory designation on TSP changes, and all other variables are as defined in Section 4.1. The EPA assigns nonattainment status to counties where TSP concentrations exceed federally mandated thresholds, subjecting those counties to substantially more stringent emissions regulations than attainment counties. This regulatory variation generates exogenous reductions in TSP that are driven by the discontinuous application of federal standards rather than by local economic

conditions, providing the variation needed to isolate the causal effect of pollution on suicide rates.

For nonattainment status to be a valid instrument, two conditions must hold. First, it must be strongly correlated with TSP changes, the relevance condition. Second, it must affect suicide rates only through its effect on air quality, the exclusion restriction. Table 5 provides evidence supporting the exclusion restriction by reporting the relationship between nonattainment status and the observable economic control variables. The estimated associations between nonattainment status and changes in income, unemployment, manufacturing employment, population density, and urban share are all small in magnitude and statistically insignificant, consistent with the instrument being orthogonal to the economic shocks that are the primary source of bias in the OLS estimates.

Table 6 reports the first-stage relationship between nonattainment status and TSP changes. The effect of nonattainment designation on TSP is large, negative, and highly significant across all three specifications, with an estimate of approximately $-10 \mu\text{g}/\text{m}^3$. The stability of this estimate across specifications that progressively add controls and interactions is consistent with the finding in Table 5 that nonattainment status is uncorrelated with observable economic conditions. Table 7 reports the reduced-form relationship between nonattainment status and suicide rate changes directly.

3.4.4 2SLS Results

Tables 8 and 9 present the two-stage least squares estimates using 1975–1976 and 1975 nonattainment status as instruments, respectively. Across both instrument definitions and both specifications reported, the estimated effect of TSP on suicide rates is negative but small in magnitude and statistically insignificant. The point estimates range from -0.025 to -0.050 , implying that a $10 \mu\text{g}/\text{m}^3$ reduction in TSP, roughly the average treatment effect documented in the first stage, would be associated with a change in the age-adjusted suicide rate of at most 0.50 per 100,000. This effect is not statistically distinguishable from zero, and it is also not large relative to the sample mean.

The 2SLS estimates are somewhat larger in magnitude when the 1975–1976 instrument is used rather than the 1975 instrument alone. This pattern is consistent with the interpretation that the regulatory effect of nonattainment status on pollution diminishes over time as more counties achieve compliance. Also, a marginal county assigned nonattainment status in later years faces a smaller required reduction. Tables 10 and 11 show robustness checks by restricting the sample to counties with 1974 TSP levels in the ranges $[50, 75] \mu\text{g}/\text{m}^3$ and $[50, 100] \mu\text{g}/\text{m}^3$ respectively. These results focus on counties near the nonattainment threshold where the instrument is most likely to generate variation. The result is null across these restricted samples, though the estimates are imprecise given the smaller sample sizes.

To think whether the null result reflects absence of an effect or insufficient precision, it is useful to consider the confidence intervals around the point estimates. The 95% confidence interval around the largest point estimate spans approximately $[-0.148, 0.048]$, implying that

a $10 \mu\text{g}/\text{m}^3$ TSP reduction could be associated with a change in the age-adjusted suicide rate anywhere from -1.48 to $+0.48$ per 100,000. The baseline age-adjusted suicide rate in this period was approximately 12 per 100,000, the upper bound of the confidence interval, so we can rule out reductions larger than 11% of the baseline rate. The estimates are more consistent with a null effect. The null result should be interpreted with some caution. While the point estimates are small and insignificant, the precision of the estimates is limited by the available sample size and the noisiness of county-level suicide data.

Table 3.16 presents population-weighted estimates, where counties are weighted by their 1970–1980 average population, giving more influence to large counties such as Los Angeles or Cook County than to small rural ones. This is the appropriate estimand if the goal is to characterize the effect experienced by the average person rather than the average county. The weighted OLS estimates are similar in magnitude to the unweighted results. The weighted 2SLS estimates differ substantially: without controls, the coefficient on TSP is positive and significant at the 1% level, implying that higher TSP concentrations are associated with higher suicide rates in large, populous counties — the direction consistent with pollution harming mental health. With controls, the weighted 2SLS estimate remains positive but loses statistical significance.

The divergence between weighted and unweighted IV estimates warrants caution. Large urban counties are more likely to have multiple monitors, making the county mean TSP a more reliable measure, but they are also subject to more severe economic shocks and spatially heterogeneous pollution exposure. Weighting them more heavily could amplify the confounding from local economic conditions that the instrument is designed to address.

There is also a first-stage concern: the regulatory response to nonattainment designation may differ systematically between large and small counties, making the instrument’s effect on TSP heterogeneous in ways that interact with population size. Given these considerations, the weighted results are best interpreted as suggestive evidence that the effect of pollution on suicide rates may be more detectable in large urban counties, rather than as a definitive causal estimate.

3.5 Heterogeneity Analysis

3.5.1 Motivation

Social, economic, and biological factors generate substantial variation in baseline suicide risk across demographic groups, and there are independent reasons to expect that both pollution exposure and vulnerability to its health effects differ systematically by race and gender. If the effect of air pollution on mental health is concentrated in particular subpopulations, an analysis of the full sample may lack the statistical power to detect it. This section explores heterogeneity by race and gender by applying the same first-difference instrumental variable approach to demographic subsamples constructed from the CDC WONDER data.

3.5.2 Racial Heterogeneity

There is well-documented evidence that the Black population bears a disproportionate share of air pollution exposure in the United States. [Currie et al. \(2020\)](#) analyze exposure to

ambient particulate matter between 2000 and 2015 and find that Black communities tend to live in the most polluted areas nationally. The mean gap in pollution exposure between African Americans and non-Hispanic whites narrowed from approximately $1.5 \mu\text{g}/\text{m}^3$ to $0.5 \mu\text{g}/\text{m}^3$ over this period. This convergence accounts for roughly 7% of the improvement in the Black-white life expectancy gap. Importantly, this racial disparity in exposure cannot be explained by differences in individual characteristics such as income and education. It suggests that residential segregation and the geographic distribution of industrial activity are the primary drivers. The authors connect the convergence to the disproportionate improvement in air quality in majority-Black areas, which they link to the enforcement of Clean Air Act regulations.

These patterns provide motivation for examining racial heterogeneity in the present setting. The heterogeneity analysis estimates the same first-difference instrumental variable specification described in Section 4.2 separately for white and non-white subsamples, using county-level TSP nonattainment status as the instrument in both cases. The instrument does not exploit within-county racial variation in pollution exposure, but it captures county-level TSP reductions driven by Clean Air Act enforcement. The motivation for examining the white subsample separately is driven by statistical power. White males have substantially higher baseline suicide rates than other demographic groups, providing greater variance in the outcome and hence a better-powered test of whether TSP reductions affected suicide mortality.

3.5.3 Gender Heterogeneity

Male suicide rates are substantially higher than female rates across all age groups. This pattern is well-established in the literature ([Daly and Wilson \(2006\)](#)) and visible in the CDC WONDER data. As shown in Figure 1, male age-adjusted suicide rates are roughly three times higher than female rates throughout the 1968–1978 period. The women subsample, with 582 observations, against 843 for male subsamples, is the smallest in the analysis, limiting precision for that group.

The age profile of suicide risk also differs by gender. For men, suicide risk is elevated among both younger and older age groups relative to middle-aged men, producing a U-shaped age profile. For women, by contrast, suicide risk increases approximately linearly with age. To address this, the age-adjusted white subsample is included with the crude rate white subsample. It allows for a comparison that controls for differences in county age distributions.

3.5.4 Results

The heterogeneity analysis applies the same first-difference instrumental variable specification to demographic subsamples:

$$(\tilde{Y}_i^{80} - \tilde{Y}_i^{70}) = \beta(\tilde{P}_i^{80} - \tilde{P}_i^{70}) + \alpha(X_i^{80} - X_i^{70}) + (\tilde{e}_i^{80} - \tilde{e}_i^{70}) \quad (3.1)$$

where $\tilde{Y}_i$ denotes the demographic-specific age-adjusted suicide rate for county i , and all other variables are as defined in Section 4.2.

Table 3.15 confirms that the instrument performs consistently across all demographic subsamples. The nonattainment coefficient ranges from -10.487 for the women subsample to -10.961 for the white age-adjusted subsample, with F-statistics ranging from 33.7 to 47.0 across subgroups. All first-stage estimates are significant at the 1% level. The instrument has sufficient strength in all subsamples.

Tables 3.12, 3.14, and 3.13 present OLS and 2SLS estimates of the effect of TSP on suicide rates across all six subgroups. The OLS estimates are uniformly small and negative, ranging from -0.005 for women to -0.082 for white males, and none are statistically significant. This pattern is consistent with the full-sample OLS result and similarly likely to reflect omitted variable bias rather than a causal relationship.

The 2SLS estimates without controls, reported in Table 3.14, are larger in magnitude than the corresponding OLS estimates for most subgroups. It is consistent with attenuation bias from measurement error in TSP being corrected by the instrument. White males show the largest point estimate at -0.155 , while women show the smallest at 0.003 . When controls are added (Table 3.13), the estimates drop closer to zero. White males estimates are falling to -0.090 and women to -0.021 . Across both specifications and all six subgroups, no estimate is statistically distinguishable from zero.

The null result is uniform across demographic subgroups and is not driven by offsetting effects across the population. White males consistently show the largest point estimates

in both OLS and 2SLS specifications; the women subsample shows the smallest and least precisely estimated effects, consistent with the lower baseline suicide rates and smaller sample size for that group. The interpretation is that TSP reductions of the magnitude generated by Clean Air Act enforcement did not have an impact on suicide rates.

The number of observations varies across demographic subgroups because CDC WONDER suppresses county-level mortality counts below 10 to protect privacy. Counties where the number of suicides within a demographic cell falls below this threshold are excluded from the relevant subsample, reducing the effective sample size for narrower demographic groups such as women.

3.6 Limitations

There are several limitations on the interpretation of the results. The most significant concerns the reliability of the outcome measure. Determining suicide intent is not a straightforward process. Medical examiners are not required to have psychological training, making misclassification a persistent concern. A proper assessment of suicide intent typically requires collecting information from family members, and the social stigma surrounding suicide creates strong incentives for underreporting. Classification practices also vary across jurisdictions during the study period. In some locations the absence of a suicide note was sufficient for a suicide classification regardless of other evidence. The CDC Compressed Mortality File represents the best available source of population-level suicide data for this period. However, systematic underreporting was likely more severe in the 1970s than in sub-

sequent decades. It implies that the true effect of air pollution on suicide mortality may be larger than the estimates suggest. A null result on measured suicide rates therefore cannot be taken as conclusive evidence of no effect on mental health more broadly.

There are also threats to the identification strategy that the available data do not allow us to address. Local economic activity affects suicide rates through multiple channels simultaneously. The construction of a new industrial plant, for instance, is likely to increase ambient pollution levels, raise local employment and income, and increase occupational stress. Each channel would be pushing suicide rates in a different direction. An alternative instrument based on wind direction would more cleanly separate pollution exposure from local economic conditions, but reliable wind direction data for the study period are not available at the county level.

Covered counties are generally skewed towards being more populous than the average US county. Regulated counties had higher population densities than unregulated counties. The results speak primarily to the effect of TSP reductions in monitored counties, which accounted for 80–85% of the US population, and may not generalize to the rural unmonitored counties excluded from the sample. The concentration of monitored counties in urban and industrial areas is a limitation of the research design, though the large population share covered suggests the estimates are relevant for the majority of Americans.

3.7 Conclusion

This paper examines the effect of the 1970 Clean Air Act on mental health, using county-level suicide rates as the outcome of interest. Following the identification strategy of [Chay and Greenstone \(2005\)](#), I exploit variation in TSP nonattainment status as an instrumental variable for changes in air pollution between 1970 and 1980, estimating first-differenced equations that absorb time-invariant county characteristics. OLS estimates have a counterintuitive coefficient on TSP across all subgroups, with the largest effect found for white males at -0.082 and the smallest for women at -0.005 . These estimates are consistent with omitted variable bias arising from the correlation between pollution levels and broader economic conditions. Two-stage least squares estimates are not statistically significant in the full sample or in any demographic subsample. Weak instrument concerns are not an explanation for the null result.

The findings are consistent with a subset of the existing literature. [Wang, Eliot, Koutrakis, Gryparis, Schwartz, Coull, Mittleman, Milberg, Lipsitz, and Wellenius \(2014\)](#) similarly find no association between ambient air pollution and depressive symptoms in an older adult population. [Bakian, Huber, Coon, Gray, Wilson, McMahon, and Renshaw \(2015\)](#) find evidence of a short-term effect of acute pollution exposure on suicide completion. The divergence between these results and the null finding reported here may reflect differences in the pollutants studied, the time horizon of exposure, the populations examined, or the identification strategies employed. In particular, this analysis captures the effect of sustained changes in ambient TSP concentrations rather than short-term fluctuations in daily pollution levels.

This paper makes a contribution to the growing literature on the environmental determinants of mental health. The question of whether air pollution affects suicide risk has direct implications for the social cost of pollution and the optimal design of environmental regulations. If mental health effects are present but are not captured by analyses of this type due to data limitations and measurement constraints, then current estimates of the health costs of pollution are likely to be underestimated.

There are several directions for future research from the limitations identified here. Access to daily or monthly county-level mortality data, combined with high-frequency pollution monitoring records, would allow for a more powerful test of the short-run relationship between pollution exposure and suicide risk. It would permit the inclusion of the temporal fixed effects that the present analysis cannot do. Expanding the outcome measures beyond suicide mortality to include hospitalizations for self-harm, emergency department visits, or survey-based measures of psychological distress would allow us to examine effects across the full distribution of mental health outcomes rather than of the most extreme. Finally, replicating this analysis for more recent periods, where both pollution monitoring and mortality data are more detailed and reliable, would be useful to the historical evidence examined here.

In future work, the Bartik instrument could be implemented ([Goldsmith-Pinkham, Sorkin, and Swift \(2020\)](#)). Using the 1970 Census of Population, I would construct county-level employment shares across industry categories capturing each county's initial industrial composition before the Clean Air Act took effect. These baseline shares would then be interacted with national employment growth rates by industry between 1970 and 1980. The predicted employment growth measure captures how much a county's economy would have grown based

only on its pre-existing industrial structure, rather than local factors. This predicted economic shock would be an additional control variable. It would absorb the variation in local economic conditions that may have simultaneously driven both TSP reductions and changes in suicide rates during the 1970s. For this approach to be valid, we make the assumption that a county's 1970 industry composition is uncorrelated with unobserved determinants of suicide rate changes over the time period. Implementing this robustness check would address the concern that economic confounding drives the results.

Table 3.1: Descriptive statistics

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Δ TSP	894	-8.59	22.02	-229.61	-19.54	3.81	64.99
Δ suicide rate	894	0.46	7.62	-51.20	-2.80	4.10	46.80

Unit of observation is a county. Δ TSP is the change in county geometric mean total suspended particulates ($\mu\text{g}/\text{m}^3$) between 1970 and 1980. Δ suicide rate is the change in county age-adjusted suicide rate per 100,000 between 1970 and 1980. TSP data from [Chay and Greenstone \(2005\)](#); suicide data from CDC WONDER Compressed Mortality File.

Table 3.2: Differences in Sample Means between Groups of Counties, Defined by TSPs Nonattainment Status

	<i>Dependent variable:</i>	
	Δ TSP ($\mu\text{g}/\text{m}^3$)	Δ Suicide Rate (per 100,000)
	(1)	(2)
TSP Nonattainment 1975–76	-10.784*** (1.82314)	0.485 (0.51)
Observations	894	894
R ²	0.049	0.001
Adjusted R ²	0.048	-0.0003

Note:

*p<0.1; **p<0.05; ***p<0.01

Unit of observation is a county. Each column reports the coefficient from a regression of the dependent variable on a binary indicator for TSP nonattainment status in 1975–76. Column (1) dependent variable is the change in county mean TSP ($\mu\text{g}/\text{m}^3$); column (2) is the change in age-adjusted suicide rate per 100,000. Both changes measured between 1970 and 1980. Robust standard errors (HC1) in parentheses.

Table 3.3: Correlation Matrix

	Δ sr	Δ TSP	Δ income.	Δ unemp	Δ mnfcg	Δ dens	Δ urban
Δ sr	1	-0.061	0.016	0.006	0.012	-0.027	-0.052
Δ TSP	-0.061	1	0.113	-0.076	0.165	0.069	0.018
Δ income	0.016	0.113	1	-0.361	0.156	0.247	0.204
Δ unemp	0.006	-0.076	-0.361	1	-0.376	-0.109	-0.091
Δ mnfcg	0.012	0.165	0.156	-0.376	1	0.034	-0.048
Δ dens	-0.027	0.069	0.247	-0.109	0.034	1	0.049
Δ urban	-0.052	0.018	0.204	-0.091	-0.048	0.049	1

Unit of observation is a county. All variables are changes between 1970 and 1980. Δ sr: age-adjusted suicide rate per 100,000; Δ tsp: mean TSP ($\mu\text{g}/\text{m}^3$); Δ income: income per capita (\$); Δ unemp: unemployment rate; Δ mnfcg: manufacturing employment share; Δ dens: population density; Δ urban: urban share.




Figure 3.1: Centers for Disease Control and Prevention, National Center for Health Statistics. National Vital Statistics System, Mortality: Compressed Mortality File 1968-1978. CDC WONDER Online Database, compiled from Compressed Mortality File CMF 1968-1988, Series 20, No. 2A, 2000.

Table 3.4: OLS Estimation of Effect of Air Pollution on Suicide Rates

	<i>Dependent variable: Δ Age-Adjusted Suicide Rate</i>		
	(1)	(2)	(3)
Δ TSP	-0.021* (0.014)	-0.023* (0.014)	-0.023* (0.014)
Controls	No	Yes	Yes
Interactions	No	No	Yes
Polynomials	No	No	Yes
Observations	894	891	891
R ²	0.004	0.029	0.125
Adjusted R ²	0.003	0.011	0.052

Note:

*p<0.1; **p<0.05; ***p<0.01

Unit of observation is a county. Dependent variable is the change in age-adjusted suicide rate per 100,000 between 1970 and 1980. Δ TSP is the change in county mean TSP ($\mu\text{g}/\text{m}^3$). Column (1) includes no controls; column (2) adds county-level economic and demographic controls; column (3) further adds polynomials and interactions. Robust standard errors (HC1) in parentheses.

Table 3.5: Differences in Sample Means between Groups of Counties, Defined by TSPs Nonattainment Status

	<i>Dependent variable:</i>				
	Income per capita	Unemployment	mnfctremp	Pop. density	urban
	(1)	(2)	(3)	(4)	(5)
TSPs Nonattainment	15.32 (45.76)	0.002 (0.002)	-0.002 (0.003)	-34.39 (28.89)	-0.002 (0.01)
Observations	894	894	894	894	894

Note:

*p<0.1; **p<0.05; ***p<0.01

Unit of observation is a county. Each column reports the coefficient from a regression of the dependent variable on a binary indicator for TSP nonattainment status in 1975–76. Dependent variables are changes in county-level economic characteristics between 1970 and 1980. mnfctremp: manufacturing employment share; urban: urban population share. Standard errors in parentheses.

Table 3.6: First-stage relationship between regulation (tsp7576) and air pollution changes

<i>Dependent variable: $\Delta\ tsp$</i>			
	(1)	(2)	(3)
TSP7576	-9.853954*** (1.544128)	-11.120710*** (1.671785)	-11.024760*** (1.740847)
Controls	No	Yes	Yes
Interactions	No	No	Yes
Polynomials	No	No	Yes
Observations	1,000	891	891
R ²	0.039206	0.097997	0.151610
Adjusted R ²	0.038243	0.080433	0.081427

Note: *p<0.1; **p<0.05; ***p<0.01

Unit of observation is a county. Dependent variable is the change in county mean TSP ($\mu\text{g}/\text{m}^3$) between 1970 and 1980. TSP7576 is a binary indicator for TSP nonattainment status in 1975–76. This is the first stage of the 2SLS estimates reported in Tables 8 and 9. Robust standard errors (HC1) in parentheses.

Table 3.7: Reduced-form relationship between regulation and suicide rates changes

<i>Dependent variable: $\Delta\ sr$</i>			
	(1)	(2)	(3)
TSP7576	0.484628 (0.566515)	0.561524 (0.600876)	0.205253 (0.613032)
Controls	No	Yes	Yes
Interactions	No	No	Yes
Polynomials	No	No	Yes
Observations	894	891	891
R ²	0.000820	0.026159	0.120752
Adjusted R ²	-0.000300	0.007196	0.048016

Note: *p<0.1; **p<0.05; ***p<0.01

Unit of observation is a county. Dependent variable is the change in age-adjusted suicide rate per 100,000 between 1970 and 1980. TSP7576 is a binary indicator for TSP nonattainment status in 1975–76. This is the reduced-form relationship between the instrument and the outcome. Robust standard errors (HC1) in parentheses.

Table 3.8: 2SLS with ts7576 as instrument

<i>Dependent variable: Δsr</i>		
	(1)	(2)
TSP7576	-0.045 (0.047)	-0.05 (0.05)
Controls	No	Yes
Interactions	No	No
Polynomials	No	No
Observations	894	870
R ²	-0.001	0.021
Adjusted R ²	-0.002	0.001

Note: *p<0.1; **p<0.05; ***p<0.01

Unit of observation is a county. Dependent variable is the change in age-adjusted suicide rate per 100,000 between 1970 and 1980. Δ TSP is instrumented by TSP nonattainment status in 1975–76. Column (1) includes no controls; column (2) adds county-level economic and demographic controls. Robust standard errors (HC1) in parentheses.

Table 3.9: 2SLS with ts75 as instrument

<i>Dependent variable: Δsr</i>		
	(1)	(2)
TSP75	-0.025 (0.046)	-0.031 (0.048)
Controls	No	Yes
Interactions	No	No
Polynomials	No	No
Observations	873	870
R ²	0.003	0.021
Adjusted R ²	0.002	0.001

Note: *p<0.1; **p<0.05; ***p<0.01

Unit of observation is a county. Dependent variable is the change in age-adjusted suicide rate per 100,000 between 1970 and 1980. Δ TSP is instrumented by TSP nonattainment status in 1975. Column (1) includes no controls; column (2) adds county-level economic and demographic controls. Robust standard errors (HC1) in parentheses.

Table 3.10: 2SLS with ts75 as instrument for counties with 1974 TSPs below $75 \mu\text{g}/\text{m}^3$ and above $50 \mu\text{g}/\text{m}^3$

<i>Dependent variable: Δsr</i>		
	(1)	(2)
TSP75	-0.011 (0.128)	-0.042 (0.144)
Controls	No	Yes
Interactions	No	No
Polynomials	No	No
Observations	376	373
R ²	0.006	0.065
Adjusted R ²	0.003	0.020

Note: *p<0.1; **p<0.05; ***p<0.01

Unit of observation is a county. Dependent variable is the change in age-adjusted suicide rate per 100,000 between 1970 and 1980. Δ TSP is instrumented by TSP nonattainment status in 1975. Sample restricted to counties with 1974 TSP levels between 50 and $75 \mu\text{g}/\text{m}^3$, focusing on counties near the nonattainment threshold. Robust standard errors (HC1) in parentheses.

Table 3.11: 2SLS with ts75 as instrument for counties with 1974 TSPs below $100 \mu\text{g}/\text{m}^3$ and above $50 \mu\text{g}/\text{m}^3$

<i>Dependent variable: Δsr</i>		
	(1)	(2)
TSP75	-0.086 (0.061)	-0.093 (0.064)
Controls	No	Yes
Interactions	No	No
Polynomials	No	No
Observations	424	421
R ²	-0.013	0.027
Adjusted R ²	-0.015	-0.014

Note: *p<0.1; **p<0.05; ***p<0.01

Unit of observation is a county. Dependent variable is the change in age-adjusted suicide rate per 100,000 between 1970 and 1980. Δ TSP is instrumented by TSP nonattainment status in 1975. Sample restricted to counties with 1974 TSP levels between 50 and $100 \mu\text{g}/\text{m}^3$, focusing on counties near the nonattainment threshold. Robust standard errors (HC1) in parentheses.

Table 3.12: OLS Heterogeneity: Effect of TSP on Suicide Rates by Subgroup

	<i>Dependent variable: Δ suicide rate</i>					
	Full	White	Over 20	White male	White adj.	Women
	(1)	(2)	(3)	(4)	(5)	(6)
Δ TSP	-0.021 (0.014)	-0.037 (0.030)	-0.033 (0.023)	-0.082 (0.063)	-0.031 (0.025)	-0.005 (0.016)
Controls	<i>No</i>	<i>No</i>	<i>No</i>	<i>No</i>	<i>No</i>	<i>No</i>
Observations	894	877	888	843	882	582
R ²	0.004	0.006	0.004	0.008	0.004	0.001
Adjusted R ²	0.003	0.005	0.003	0.007	0.003	-0.001

Note: *p<0.1; **p<0.05; ***p<0.01. Robust standard errors (HC1) in parentheses. Unit of observation is a county. Each column reports OLS estimates for a different demographic subsample. Dependent variable is the change in the subgroup-specific age-adjusted suicide rate per 100,000 between 1970 and 1980, from CDC WONDER Compressed Mortality File. Δ TSP is the change in county mean TSP ($\mu\text{g}/\text{m}^3$). No controls included. Subgroup samples are smaller than the full sample due to CDC WONDER suppression of counts below 10.

Table 3.13: 2SLS Heterogeneity: Effect of TSP on Suicide Rates by Subgroup, With Controls

	<i>Dependent variable: Δ suicide rate</i>					
	Full	White	Over 20	White male	White adj.	Women
	(1)	(2)	(3)	(4)	(5)	(6)
Δ TSP	-0.050 (0.050)	-0.097 (0.090)	-0.048 (0.049)	-0.090 (0.107)	-0.068 (0.073)	-0.021 (0.052)
Controls	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>
Instrument	TSP Nonattainment 1975–76					
Observations	891	874	885	840	879	579
R ²	0.024	0.013	0.032	0.026	0.017	0.075
Adjusted R ²	0.005	-0.005	0.013	0.007	-0.001	0.047

Note: *p<0.1; **p<0.05; ***p<0.01. Robust standard errors (HC1) in parentheses. Unit of observation is a county. Each column reports 2SLS estimates for a different demographic subsample, with county-level economic and demographic controls. Dependent variable is the change in the subgroup-specific age-adjusted suicide rate per 100,000 between 1970 and 1980, from CDC WONDER. Δ TSP is instrumented by TSP nonattainment status in 1975–76. Subgroup samples are smaller than the full sample due to CDC WONDER suppression of counts below 10.

Table 3.14: 2SLS Heterogeneity: Effect of TSP on Suicide Rates by Subgroup, No Controls

	<i>Dependent variable: Δ suicide rate</i>					
	Full	White	Over 20	White male	White adj.	Women
	(1)	(2)	(3)	(4)	(5)	(6)
Δ TSP	-0.045 (0.047)	-0.100 (0.088)	-0.042 (0.047)	-0.155 (0.133)	-0.075 (0.073)	0.003 (0.050)
Controls	<i>No</i>	<i>No</i>	<i>No</i>	<i>No</i>	<i>No</i>	<i>No</i>
Instrument	TSP Nonattainment 1975–76					
Observations	894	877	888	843	882	582
R ²	-0.001	-0.011	0.004	0.002	-0.005	-0.001
Adjusted R ²	-0.002	-0.012	0.003	0.0004	-0.006	-0.002

Note: *p<0.1; **p<0.05; ***p<0.01. Robust standard errors (HC1) in parentheses. Unit of observation is a county. Each column reports 2SLS estimates for a different demographic subsample, with no controls. Dependent variable is the change in the subgroup-specific age-adjusted suicide rate per 100,000 between 1970 and 1980, from CDC WONDER. Δ TSP is instrumented by TSP nonattainment status in 1975–76. Subgroup samples are smaller than the full sample due to CDC WONDER suppression of counts below 10.

Table 3.15: First Stage by Subgroup Sample

	<i>Dependent variable: Δ TSP</i>					
	Full	White	Over 20	White male	White adj.	Women
	(1)	(2)	(3)	(4)	(5)	(6)
TSP Nonattainment 1975–76	-10.784*** (1.597)	-10.829*** (1.599)	-10.939*** (1.596)	-10.735*** (1.620)	-10.961*** (1.599)	-10.487*** (1.807)
Observations	894	877	888	843	882	582
R ²	0.049	0.050	0.050	0.050	0.051	0.055
Adjusted R ²	0.048	0.049	0.049	0.049	0.050	0.053
F-statistic	45.6***	45.9***	47.0***	43.9***	47.0***	33.7***

Note: *p<0.1; **p<0.05; ***p<0.01. Robust standard errors (HC1) in parentheses. Unit of observation is a county. Dependent variable is the change in county mean TSP ($\mu\text{g}/\text{m}^3$) between 1970 and 1980. TSP Nonattainment 1975–76 is a binary indicator for TSP nonattainment status in 1975–76. Each column reports the first stage for the corresponding demographic subsample used in Tables 14 and 15. F-statistics test the null of instrument irrelevance.

Table 3.16: Robustness Check: Population-Weighted Estimates

	<i>Dependent variable: Δsr</i>			
	OLS		2SLS	
	(1)	(2)	(3)	(4)
ΔTSP	-0.021** (0.008)	-0.023*** (0.009)	0.115*** (0.039)	0.055 (0.042)
Controls	No	Yes	No	Yes
Instrument	—	—	TSP7576	TSP7576
Observations	894	891	894	873
R ²	0.007	0.114	-0.292	0.031
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01. Population weights used.			

References

- ARCHSMITH, J., A. HEYES, AND S. SABERIAN (2018): “Air quality and error quantity: Pollution and performance in a high-skilled, quality-focused occupation,” *Journal of the Association of Environmental and Resource Economists*, 5, 827–863.
- BAKIAN, A. V., R. S. HUBER, H. COON, D. GRAY, P. WILSON, W. M. MCMAHON, AND P. F. RENSHAW (2015): “Acute air pollution exposure and risk of suicide completion,” *American journal of epidemiology*, 181, 295–303.
- BLINCOE, L., T. MILLER, J.-S. WANG, D. SWEDLER, T. COUGHLIN, B. LAWRENCE, F. GUO, S. KLAUER, AND T. DINGUS (2023): “The Economic and Societal Impact of Motor Vehicle Crashes, 2019 (Revised),” .
- BLOCK, M. AND L. CALDERON-GARCIDUENAS (2009): “Air Pollution: Mechanisms of Neuroinflammation CNS Disease,” *Trends in neurosciences*, 32, 506–16.
- BRODEUR, A., N. COOK, AND T. WRIGHT (2021): “On the Effects of COVID-19 Safer-At-Home Policies on Social Distancing, Car Crashes and Pollution,” GLO Discussion Paper Series 774, Global Labor Organization (GLO).
- BROOK, R. D., S. RAJAGOPALAN, C. A. POPE III, J. R. BROOK, A. BHATNAGAR, A. V. DIEZ-ROUX, F. HOLGUIN, Y. HONG, R. V. LUEPKER, M. A. MITTLEMAN, ET AL. (2010): “Particulate matter air pollution and cardiovascular disease: An update to the scientific statement from the American Heart Association,” *Circulation*, 121, 2331–2378.
- BURTON, A. M. AND T. ROACH (2026): “Negative externalities of temporary reductions in

- cognition: Evidence from particulate matter pollution and fatal car crashes,” *Journal of Environmental Economics and Management*, 138, 103310.
- CALDERON-GARCIDUENAS, L., M. FRANCO-LIRA, R. TORRES-JARDON, C. HENRIQUEZ-ROLDAN, G. BARRAGAN-MEJIA, G. VALENCIA-SALAZAR, A. GONZALES-MACIEL, R. REYNOSO-ROBLES, A. VILLARREAL-CALDERON, AND W. REED (2008): “Air pollution, cognitive deficits and brain abnormalities: a pilot study with children and dogs,” *Brain and cognition*, 68, 117–127.
- CARSON, R. T., P. KOUNDOURI, AND C. NAUGES (2011): “Arsenic mitigation in Bangladesh: a household labor market approach,” *American Journal of Agricultural Economics*, 93, 407–414.
- CHAY, K. Y. AND M. GREENSTONE (2005): “Does air quality matter? Evidence from the housing market,” *Journal of political Economy*, 113, 376–424.
- COHEN, A. AND L. EINAV (2003): “The effects of mandatory seat belt laws on driving behavior and traffic fatalities,” *Review of Economics and Statistics*, 85, 828–843.
- CONGRESSIONAL BUDGET OFFICE (2012): “An Update to the Budget and Economic Outlook: Fiscal Years 2012 to 2022,” Tech. rep., Congressional Budget Office, Washington, DC.
- CURRIE, J., J. VOORHEIS, AND R. WALKER (2020): “What caused racial disparities in particulate exposure to fall? New evidence from the Clean Air Act and satellite-based measures of air quality,” Tech. rep., National Bureau of Economic Research.
- DALY, M. C. AND D. J. WILSON (2006): “Keeping up with the Joneses and staying ahead of the Smiths: evidence from suicide data,” *FRB of San Francisco Working Paper*.

- DANAGOULIAN, S. AND M. DEZA (2024): “Driving Under the Influence of Allergies: The Effect of Seasonal Pollen on Traffic Fatalities,” Tech. rep., National Bureau of Economic Research.
- DERYUGINA, T., G. HEUTEL, N. H. MILLER, D. MOLITOR, AND J. REIF (2019): “The mortality and medical costs of air pollution: Evidence from changes in wind direction,” *American Economic Review*, 109, 4178–4219.
- DOCKERY, D. W., C. A. POPE, X. XU, J. D. SPENGLER, J. H. WARE, M. E. FAY, B. G. FERRIS, AND F. E. SPEIZER (1993): “An association between air pollution and mortality in six US cities,” *New England Journal of Medicine*, 329, 1753–1759.
- EISENBERG, D. (2003): “Evaluating the effectiveness of policies related to drunk driving,” *Journal of policy Analysis and management*, 22, 249–274.
- EPA, U. (2006): “Report of the EPA work group on VSL meta-analysis. NCEE-0494,” *National Center for Environmental Economics, EPA, Washington*.
- FIORDELISI, A., P. PISCITELLI, B. TRIMARCO, E. COSCIONI, G. IACCARINO, AND D. SORRIENTO (2017): “The mechanisms of air pollution and particulate matter in cardiovascular diseases,” *Heart failure reviews*, 22, 337–347.
- GIULIETTI, C., M. TONIN, AND M. VLASSOPOULOS (2020): “When the market drives you crazy: Stock market returns and fatal car accidents,” *Journal of Health Economics*, 70, 102245.
- GO, T.-H., M.-H. KIM, Y.-Y. CHOI, J. HAN, AND C. KIM (2024): “The short-term effect of ambient particulate matter on suicide death,” *Environmental Health*, 23.

- GOLDSMITH-PINKHAM, P., I. SORKIN, AND H. SWIFT (2020): “Bartik Instruments: What, When, Why, and How,” *American Economic Review*, 110, 2586–2624.
- GREENBERG, P. E., A.-A. FOURNIER, T. SISITSKY, C. T. PIKE, AND R. C. KESSLER (2015): “The economic burden of adults with major depressive disorder in the United States (2005 and 2010),” *Journal of Clinical Psychiatry*, 76, 155–162.
- GUPTA, N., M. U. M. JOHARI, H. JASHAMI, AND P. T. SAVOLAINEN (2022): “How is traffic safety affected by changes in traffic speeds following speed limit increases? An evaluation with probe vehicle data,” *Traffic Safety Research*, 3, 000017.
- HANNA, R. AND P. OLIVA (2015): “The effect of pollution on labor supply: Evidence from a natural experiment in Mexico City,” *Journal of Public Economics*, 122, 68–79.
- HAUSMAN, J. A., B. D. OSTRO, AND D. A. WISE (1984): “Air pollution and lost work,” .
- HE, J., H. LIU, AND A. SALVO (2019): “Severe air pollution and labor productivity: Evidence from industrial towns in China,” *American Economic Journal: Applied Economics*, 11, 173–201.
- HERRNSTADT, E., A. HEYES, E. MUEHLEGGGER, AND S. SABERIAN (2021): “Air pollution and criminal activity: Microgeographic evidence from Chicago,” *American Economic Journal: Applied Economics*, 13, 70–100.
- HOEK, G., R. M. KRISHNAN, R. BEELEN, A. PETERS, B. OSTRO, B. BRUNEKREEF, AND J. D. KAUFMAN (2013): “Long-term air pollution exposure and cardio-respiratory mortality: a review,” *Environmental Health*, 12, 1–15.
- ISEN, A., M. ROSSIN-SLATER, AND W. R. WALKER (2017): “Every breath you take—every

- dollar you'll make: The long-term consequences of the Clean Air Act of 1970," *Journal of Political Economy*, 125, 848–902.
- JANS, J., P. JOHANSSON, AND J. P. NILSSON (2018): "Economic status, air quality, and child health: Evidence from inversion episodes," *Journal of health economics*, 61, 220–232.
- KIM, K.-H., S. A. JAHAN, AND E. KABIR (2013): "A review on human health perspective of air pollution with respect to allergies and asthma," *Environment international*, 59, 41–52.
- KNAPP, K. K., D. KROEGER, K. GIESE, ET AL. (2000): "Mobility and safety impacts of winter storm events in a freeway environment," Tech. rep., Iowa State University. Center for Transportation Research and Education.
- LEHTINEN, V. AND M. JOUKAMAA (1994): "Epidemiology of depression: prevalence, risk factors and treatment situation," *Acta Psychiatrica Scandinavica*, 89, 7–10.
- LELIEVELD, J., J. S. EVANS, M. FNAIS, D. GIANNADAKI, AND A. POZZER (2015): "The contribution of outdoor air pollution sources to premature mortality on a global scale," *Nature*, 525, 367–371.
- LIN, G.-Z., L. LI, Y.-F. SONG, Y.-X. ZHOU, S.-Q. SHEN, AND C.-Q. OU (2016): "The impact of ambient air pollution on suicide mortality: A case-crossover study in Guangzhou, China," *Environmental Health*, 15, 90.
- LUNDBERG, P., E. CANTOR-GRAAE, G. RUKUNDO, S. ASHABA, AND P.-O. ÖSTERGREN (2009): "Urbanicity of place of birth and symptoms of psychosis, depression and anxiety in Uganda," *The British Journal of Psychiatry*, 195, 156–162.

- MCKENZIE, B. (2013): “County-to-County Commuting Flows: 2006–2010,” American community survey reports, U.S. Census Bureau, Washington, DC.
- POPE, C. A., R. T. BURNETT, M. J. THUN, E. E. CALLE, D. KREWSKI, K. ITO, AND G. D. THURSTON (2002): “Lung cancer, cardiopulmonary mortality, and long-term exposure to fine particulate air pollution,” *Jama*, 287, 1132–1141.
- POPE, C. A. AND D. W. DOCKERY (2006): “Health effects of fine particulate air pollution: lines that connect,” *Journal of the air & waste management association*, 56, 709–742.
- ROCKETT, I. R., N. D. KAPUSTA, R. BHANDARI, ET AL. (2011): “Suicide misclassification in an international context: revisitation and update,” *Suicidology online*, 2, 48–61.
- RODRIGUES-SILVA, F., U. DE PAULA SANTOS, P. H. N. SALDIVA, L. F. AMATO-LOURENÇO, S. G. E. K. MIRAGLIA, ET AL. (2012): “Health risks and economic costs of absenteeism due to air pollution in São Paulo, Brazil,” *Aerosol and Air Quality Research*, 12, 826–833.
- RUO, B., J. S. RUMSFELD, M. A. HLATKY, H. LIU, W. S. BROWNER, AND M. A. WHOOLEY (2003): “Depressive symptoms and health-related quality of life: the Heart and Soul Study,” *Jama*, 290, 215–221.
- SAGER, L. (2019): “Estimating the effect of air pollution on road safety using atmospheric temperature inversions,” *Journal of Environmental Economics and Management*, 98, 102250.
- SCHLENKER, W. AND M. J. ROBERTS (2009): “Nonlinear temperature effects indicate severe damages to US crop yields under climate change,” *Proceedings of the National Academy of Sciences*, 106, 15594–15598.

- SEGALOWITZ, S. J. (2008): “Public health, brain health, and the dangers of air pollution for neural development.” .
- SULLIVAN, D. M., A. KRUPNICK, ET AL. (2018): “Using satellite data to fill the gaps in the US air pollution monitoring network,” *Resources for the future working paper*, 18.
- SZYSZKOWICZ, M., J. B. WILLEY, E. GRAFSTEIN, B. H. ROWE, AND I. COLMAN (2010): “Air pollution and emergency department visits for suicide attempts in Vancouver, Canada,” *Environmental health insights*, 4, EHI-S5662.
- TROUBAT, R., P. BARONE, S. LEMAN, ET AL. (2021): “Neuroinflammation and depression: A review,” *European Journal of Neuroscience*, 53, 151–171.
- U.S. ENVIRONMENTAL PROTECTION AGENCY (1999): “The Benefits and Costs of the Clean Air Act 1990 to 2010: EPA Report to Congress,” Tech. Rep. EPA-410-R-99-001, Office of Air and Radiation and Office of Policy, U.S. Environmental Protection Agency, Washington, DC.
- VOS, T., C. ALLEN, M. ARORA, R. M. BARBER, Z. A. BHUTTA, A. BROWN, A. CARTER, D. C. CASEY, F. J. CHARLSON, A. Z. CHEN, ET AL. (2016): “Global, regional, and national incidence, prevalence, and years lived with disability for 310 diseases and injuries, 1990–2015: a systematic analysis for the Global Burden of Disease Study 2015,” *The lancet*, 388, 1545–1602.
- WANG, Y., M. N. ELIOT, P. KOUTRAKIS, A. GRYPARIS, J. D. SCHWARTZ, B. A. COULL, M. A. MITTLEMAN, W. P. MILBERG, L. A. LIPSITZ, AND G. A. WELLENIUS (2014): “Ambient air pollution and depressive symptoms in older adults: results from the MOBILIZE Boston study,” *Environmental health perspectives*, 122, 553–558.

WEUVE, J., R. C. PUETT, J. SCHWARTZ, J. D. YANOSKY, F. LADEN, AND F. GRODSTEIN

(2012): “Exposure to particulate air pollution and cognitive decline in older women,”

Archives of internal medicine, 172, 219–227.

WORLD BANK (2016): “The Cost of Air Pollution: Strengthening the Economic Case for

Action,” Tech. rep., World Bank Group.

WORLD HEALTH ORGANIZATION (2025): “World Mental Health Today: Latest Data,” .