

Need a dissertation on a similar topic? [Check out](#) how our dissertation services can help you.



Graduate Theses, Dissertations, and Problem Reports

2024

Three Essays on Misinformation, Mistrust and their Influence on Public Health Policy

Eli Kochersperger
West Virginia University

Follow this and additional works at: <https://researchrepository.wvu.edu/etd>



Part of the [Health Economics Commons](#)

Recommended Citation

Kochersperger, Eli, "Three Essays on Misinformation, Mistrust and their Influence on Public Health Policy" (2024). *Graduate Theses, Dissertations, and Problem Reports*. 12369.
<https://researchrepository.wvu.edu/etd/12369>

This Dissertation is protected by copyright and/or related rights. It has been brought to you by the The Research Repository @ WVU with permission from the rights-holder(s). You are free to use this Dissertation in any way that is permitted by the copyright and related rights legislation that applies to your use. For other uses you must obtain permission from the rights-holder(s) directly, unless additional rights are indicated by a Creative Commons license in the record and/ or on the work itself. This Dissertation has been accepted for inclusion in WVU Graduate Theses, Dissertations, and Problem Reports collection by an authorized administrator of The Research Repository @ WVU. For more information, please contact researchrepository@mail.wvu.edu.

Three Essays on Misinformation, Mistrust and their Influence on Public Health Policy

Eli Kochersperger

Dissertation submitted

to the John Chambers College of Business and Economics

at West Virginia University

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in

Economics

Daniel Grossman, Ph.D., Chair

Jane Ruseski, Ph.D.

Alicia Plemmons, Ph.D.

Heather Stephens, Ph.D.

Department of Economics

Morgantown, West Virginia.

2024

Keywords: Health Behavior; Belief; Unawareness

©2024 Eli Kochersperger

Abstract

Three Essays on Misinformation, Mistrust and their Influence on Public Health Policy

Eli Kochersperger

In the first chapter I identify the impact of fentanyl exposure misinformation- namely, the erroneous belief that momentary, passive contact with the potent opioid fentanyl can be seriously harmful- on first responder behavior during overdose events, and on overall opioid-related mortality. I examine changes in opioid-related mortality following one particularly well-covered episode involving an Ohio police officer in 2017, wherein the officer appeared to experience an acute opioid overdose after touching what was believed to be fentanyl. Employing a synthetic differences-in-differences identification strategy, I find areas with greater media exposure to this misinformation exhibit marked increases in opioid overdose deaths; as well as preliminary evidence to suggest that overdose interventions performed by first responders were less effective and argue that this represents an increased hesitancy to render aid due to the potential presence of fentanyl. These results point to the existence of a heretofore unrecognized driving factor behind the current opioid epidemic, as well as to the need for policy intervention to counteract further dissemination of such adverse misinformation.

The second chapter builds on these initial findings from the preceding chapter and principally examines the effects of misinformation on the actions of first responders in responding to opioid overdoses in New York State. By utilizing data that distinguishes naloxone administrations across law enforcement officers, emergency medical services, and community opioid overdose programs groups, I examine the relative change in first responder behavior by type following a well-covered media event credited with popularizing the erroneous belief that momentary contact with the opioid fentanyl is lethal. I find evidence to suggest that law enforcement officer resuscitations using naloxone following this event decreased significantly in frequency, and argue that this represents an increased hesitancy to render aid due to unfounded fears of fentanyl exposure. These results highlight the potential adverse effects that misinformation pose to public health outcomes, and offer an alternative perspective to understand previous empirical research which has shown inconsistent results regarding the effectiveness of harm reduction policies.

In the third chapter I employ microdata from the Census Bureau's experimental Household Pulse Survey to examine Medicaid participation impacts on the COVID-19 vaccination decision and find that program recipients exhibited lower proclivities to be vaccinated relative to demographically-similar households insured through some alternative provider. Additionally, I observe that administrative burden, measured by state-level easements in the program enrollment process, is associated with significant declines in self-reported vaccination hesitancy. These results suggest the psychic and time costs accompanying the enrollment or benefits utilization processes may be negatively influencing other health behaviors, and highlight the significant policy implications of loosening Medicaid enrollment protocols on public health outcomes.

Acknowledgements

There are too many people who have helped me through the process of composing this dissertation for me to meaningfully express my gratitude in a few short words, but I will nonetheless try.

First, I want to thank my advisor, Dr. Daniel Grossman, for his patience and guidance in helping this research become something actually legible, if even sensible (at least in some parts). I also want to state how grateful I am to my entire committee for their assistance—their input has demonstrated a level of professionalism and expertise that I will strive to emulate throughout my research career.

I would also like to thank the members of my cohort, Shaun Gilyard, Dinushka Parnavitana, Bryan Khoo, Justin Heflin, and Corey Williams for their friendship and support throughout grad school; my family for teaching me the importance of asking questions, and the value of good scholarship; and my wife, Kasia, for everything else. I do not need any sophisticated estimators or identification approaches to know that the causal factor which turned me from a college dropout to academic was her presence in my life.

Contents

1	Quantifying the Effects of Fentanyl Exposure Misinformation on Opioid Mortality	1
1.1	Introduction	1
1.2	Background	4
1.2.1	Fentanyl, naloxone, and first responder overdose protocols	4
1.2.2	The fentanyl misinformation panic and its origins	8
1.3	Empirical Approach	11
1.3.1	Data	11
1.3.2	Identification Strategy	17
1.4	Results	21
1.4.1	Primary Results	21
1.4.2	Robustness Checks	28
1.5	Mechanisms Analysis	34
1.5.1	Identifying media’s direct role in misinformation shocks	34
1.5.2	Identifying changes to first responder behavior	38
1.6	Conclusions	43
1.7	References	46
2	Fentanyl on My Mind: Perceived Opioid Exposure Risk and its Influence on Naloxone Administration Rates	52
2.1	Introduction	52
2.2	Background	55
2.2.1	Naloxone and New York state	55
2.2.2	The fentanyl misinformation panic	58
2.3	Empirical approach	61
2.3.1	Data	61
2.3.2	Identification strategy	63
2.4	Results	66
2.4.1	Primary results	66
2.4.2	Addressing threats to validity	71
2.5	Discussion	76
2.6	Conclusions	84
2.7	References	86
2.8	Appendices	90
2.8.1	Further mortality analysis, controlling for overdose counts	90
2.8.2	Additional web-based exposure results	92
3	Insurance Barrier Impacts on Vaccine Hesitancy: Administrative Burden and COVID-19 Vaccination within the Medicaid Population	95
3.1	Introduction	95
3.2	Background	98
3.3	Empirical strategy	102
3.3.1	Data	102
3.3.2	Methods	106
3.4	Results	111

3.4.1	Income eligibility cap analysis	118
3.4.2	Mechanisms	126
3.5	Discussion & Conclusions	130
3.6	References	134
3.7	Appendices	139
3.7.1	Propensity score matched model results	139

List of Figures

1.1	Time series of Google search interest in the hazards of fentanyl exposure.	12
1.2	Regional time series of Google search interest in fentanyl.	12
1.3	Time series for opioid mortality and prevalence.	18
1.4	Treatment maps for East Liverpool event.	24
1.5	Plotted primary SDiD results.	25
1.6	Distribution of placebo test results.	28
1.7	Comparison of SDiD controls constructed under different donor set specifications.	31
1.8	Regional time series of local and county-level government expenditures on police protection, fire prevention, and health services.	36
1.9	Regional time series of naloxone kit distribution, trainings, and administrations.	40
2.1	Average opioid-related death rates and counts timeseries county-quarter and opioid type.	64
2.2	Average naloxone administration counts timeseries by county-quarter and first responder type.	67
2.3	Event studies for naloxone administrations, by treatment definition. . . .	77
2.4	County-Level Heterogenous Treatment Effects.	84
2.5	Time series of Google search interest in the hazards of fentanyl exposure.	94
2.6	Average LEO naloxone administrations, observed and counterfactual. . .	94
3.1	County-level primary vaccination series completion rate time series, stratified by Medicaid enrollment rates.	97
3.2	Self-reported vaccination or intent to vaccinate rates time series by insurance coverage type.	98
3.3	Entropy balanced results for Medicaid recipients against uninsured subsample, stratified by percent of Medicaid income eligibility cap.	121
3.4	Regression discontinuity analysis on propensity score matched samples (Medicaid recipients against uninsured).	124
3.5	Estimated State \times Medicaid coefficients against Administrative Burden and Benefits & Copay Indexes.	127

List of Tables

1.1	Data description, coverage and sources	15
1.2	County-quarter summary statistics table.	16
1.3	SDiD coefficient estimates for mortality by drug type, stratified by demographics and age.	23
1.4	Ohio fentanyl and heroin prevalence analysis results.	27
1.5	SDiD robustness checks, DMAs as treatment unit	30
1.6	SDiD robustness checks, commuting zones as treatment unit	33
1.7	Covariate balance tables for SDiD donor set specifications.	35
1.8	Falsification test SDiD results.	35
1.9	Staggered SDiD estimates for all US counties treated with media coverage of East Liverpool event.	38
1.10	Place-of-death linear probability results.	42
2.1	Summary Statistics of New York Data by County-Quarter	62
2.2	Naloxone administration regression results, LEO as treated.	69
2.3	Naloxone administration regression results, LEO and EMS as treated.	70
2.4	Naloxone administration robustness checks results	72
2.5	Alternative continuous treatment specifications to address threats to validity	74
2.6	Opioid hospitalization regression results	79
2.7	Opioid mortality regression results	81
2.8	Mortality results, controlling for overdose counts.	91
3.1	Summary statistics for cleaned Pulse Survey data, by insurance coverage type.	105
3.2	Naive OLS estimates for likelihood of receiving or intention to receive vaccine by insurance coverage type.	112
3.3	Entropy balanced estimates for likelihood of receiving or intention to receive vaccine by insurance coverage type.	114
3.4	Entropy balanced estimates for likelihood of receiving or intention to receive vaccine by insurance coverage type, stratified by individual characteristics.	116
3.5	State stratification entropy balancing analysis by political and expansion policy status.	117
3.6	Entropy balanced likelihood estimates for vaccine hesitancy reason by insurance coverage type.	119
3.7	Regression discontinuity results on propensity score matched samples, percent of Medicaid income eligibility cap as running variable.	123
3.8	2SLS estimates using percent of income cap as instrument.	125
3.9	Entropy balanced estimates for likelihood of receiving or intention to receive vaccine by Medicaid status and state-level Medicaid measures interactions (Medicaid recipients against privately insured).	129
3.10	State stratification entropy balancing analysis by Medicaid enrollment ease.	131
3.11	Propensity score matched estimates for likelihood of receiving or intention to receive vaccine by insurance coverage type.	140
3.12	Entropy balanced estimates for likelihood of receiving or intention to receive vaccine by Medicaid status and state-level Medicaid measures interactions (Medicaid recipients against uninsured).	141

1 Quantifying the Effects of Fentanyl Exposure Misinformation on Opioid Mortality

By causing fear and panic among such key partners in responding to the overdose crisis, we're putting people's lives at further risk and adding to the stigma around drug use.

— Keith Brown, Katal Center for Health, Equity and Justice director,
quoted in the *Times-Union* in 2018.

1.1 Introduction

In 2021, life expectancy at birth in the United States declined by 1.16%, from 77.0 to 76.1 years (Arias, Tejada-Vera, Ahmad, & Kochanek, 2022). These stark changes are overwhelmingly the result of excess COVID-19 deaths, but deviations from long-term trends can be traced back, at least in part, to shifts that first presented in the years preceding the pandemic. In the years 2015 and 2017, on the heels of a worsening opioid epidemic, life expectancy reversed a 25-year trend of year-over-year gains to decline by 0.2 and 0.1 years (Devitt, 2018), respectively; and for as much attention as the COVID-19 pandemic has duly received, it is worth noting the sobering statistic that it was in 2021 wherein annual drug-involved overdose deaths exceeded 100,000 persons for the first time. For perspective, of the 0.9 year decline in life expectancies in that year, approximately 16% (compared to COVID-19's 50%) can be directly attributed to changes in unintentional injury deaths, of which accidental drug poisonings largely constitute. Within just the 18-45 age group these effects are even more pronounced, where accidental overdoses involving synthetic narcotics (an estimated 90% of which are associated with the opioid fentanyl or its analogs) were the leading single cause of death, exceeding even suicide, car accidents, heart disease and cancer (Jones, 2023).

It is unsurprising in light of these statistics that many advocacy groups and policymakers have taken to treating fentanyl as a singularly unique threat to public health. From a

pharmacological perspective this appears warranted: Fentanyl can be 100 times more potent than morphine, and 50 times more potent than heroin (Ramos-Matos, Bistas, & Lopez-Ojeda, 2022). Considering the relative ease of access and low cost of the drug in illicit markets, fentanyl is also commonly used as an adulterant which has further exacerbated its lethality through the consumption by unwitting- and often opioid naive-users. Only trace, bordering on imperceptible quantities (2 mg) of the raw substance are needed to trigger fatal respiratory failure when used intravenously, which has prompted some advocacy groups to claim that “if you can see it, it can kill you.”¹ However, like other synthetically produced opioids tramadol and carfentanil, fentanyl poses no significant acute health risk to individuals when exposed to the substance incidentally. Dermal contact and inhalation both require extremely prolonged exposure intervals to receive even clinical dosages (Moss et al., 2018), which all but eliminates the possibility of overdose from momentary contact. Nonetheless, sensational news stories detailing the supposedly-instantaneous lethality of the opioid, perhaps due in part to the embellished exposure risks promulgated by the DEA and other law enforcement agencies, have flourished (Beletsky et al., 2020). So persistent a media phenomenon has this become in fact, that public perceptions of fentanyl exposure hazards no longer align with the reality described by the clinical toxicology literature; and mere speculation on the presence of the narcotic within communities has elicited such outsized alarm as to be described as a form of moral panic (Ciccarone MD, 2020).

The question this study addresses is how these erroneous beliefs on fentanyl exposure hazards, in their near-ubiquity, have factored into the broader opioid epidemic and related overdose mortality. Previous research has succeeded in establishing the direct effect that fentanyl’s introduction to illicit drug markets has had on mortality within the context of use behavior, but has largely neglected the potential influence of bystander and first responder perceptions. This is relevant because opioid overdoses are unique among accidental drug poisonings for their relative treatability with prompt medical intervention. The opioid antagonist naloxone (also known by the brand name Narcan), can safely

¹Quote pulled from Jackson County, Missouri’s Community Backed Anti-Crime Tax (COMBAT) program website: <https://www.jacksoncountycombat.com/818/Get-The-Fentanyl-Facts>.

resuscitate unresponsive victims, requires no specialty medical training to ensure its correct administration, and is widely available without prescription at low or zero costs. Critical however, is that these naloxone interventions require close proximity between overdose victims and those rendering aid. If a first responder or bystander had the means to save an overdose victim, but also incorrectly believed that they would be at personal risk of injury in doing so, any ensuing hesitancy could easily translate to death.

Drawing on restricted-use mortality data from the National Vital Statistics System (NVSS) for the years 2014-2019, this study examines the influence of misinformation shocks on opioid-related mortality through the most common dissemination medium: Media reports of claimed “near-death” experiences suffered by first responders involving fentanyl. In particular, I focus on a 2017 incident involving an undercover narcotics officer in East Liverpool, Ohio, who was hospitalized following brief exposure to what was believed to be fentanyl powder. Prior media analyses (Beletsky et al., 2020) have suggested that this was the seminal event in pressing the fentanyl exposure myth into the public imagination. As such, I exploit the unexpected proliferation and spatial variation in media coverage of this event to estimate the association between misinformation dissemination and county-quarter opioid-related mortality rates. I find that within the media market local to East Liverpool, opioid related mortality increased significantly following the 2017 event when compared against bordering counties, and that these variations in mortality cannot be attributed to other contemporaneous factors. Similar, though slightly attenuated effects are observed within other media markets across the country that featured reporting on the East Liverpool event. Moreover, I find evidence that regions with greater media coverage of this event demonstrate marked shifts in recorded death locations- away from hospitals, and towards other third locations- among opioid overdose victims, and argue that this is the result of a reluctance to render aid based on fentanyl hazards perceptions.

The principal contribution of this paper is in providing the first credibly-causal estimates for the economic consequences of the fentanyl hazards myth, as well as more generally advancing the literature on the role of misinformation in public health policy efficacy. Focusing on mortality specifically here is essential because it directly reveals the life-

threatening impact of the epidemic and guides effective interventions to save lives. Although there is a rich body of qualitative work examining the potential influence of this misinformation on first responder behavior (Attaway, Smiley-McDonald, Davidson, & Kral, 2021; Beletsky et al., 2020; Del Pozo et al., 2021; Herman et al., 2020), with the exception of this study’s companion article (Kochersperger, 2023), no research to date has examined the direct outcomes of such beliefs on public health outcomes. Understanding the broader influence of misinformation- especially that spread through social media- on public health outcomes has been of particular attention of late on the heels of a growing anti-vaccination movement (Chou, Oh, & Klein, 2018; Wang, McKee, Torbica, & Stuckler, 2019). Within just the economics literature, for instance, Carrieri et al. (2019) employ a research design similar to that used here involving one significant media-driven misinformation shock asserting a causal relationship between receipt of the MMR vaccine and autism diagnoses to derive estimates for its effect on vaccination rates. This study therefore bridges these two literatures and bolsters the descriptive results already generated on fentanyl hazards myth phenomenon with the application of causal inference methods. The remainder of the paper is organized as follows. In the next section I offer some context on fentanyl hazards misinformation phenomena, including a summary of common first responder overdose protocols and origins of the myth to get at possible mechanisms. In Section 3 I describe my data and empirical strategy, and report my results in Section 4. I conclude with a discussion of these results and policy ramifications in Section 6.

1.2 Background

1.2.1 Fentanyl, naloxone, and first responder overdose protocols

Fentanyl is a synthetic piperidine-based opioid drug, meaning that unlike natural or semi-synthetic opioids such as morphine, heroin, or oxycodone, it is not derived from poppies. It was first developed by Paul Janssen in 1959 as an effort to create what was then the most potent analgesic, believing that to do so would improve safety (Stanley, 1992). The drug first received US medical approval in 1968 and- in line with its considerable potency-

has maintained a somewhat niche prescribing status when compared to other opioids. Primarily used for managing major pain, most fentanyl is prescribed to patients following surgery or during late-stage cancer. Among these patients it is particularly common to prescribe transdermal patches, which are adhesive strips that cling to a person's skin and are specially formulated to allow fentanyl to enter the bloodstream over prolonged periods. Beyond these intended therapeutic uses, fentanyl has a complicated and deadly legacy. The fentanyl analogs carfentanil and remifentanil were implicated in the direct deaths of 125 hostages during the 2002 Nord-Ost siege, when Russian special forces piped aerosolized forms of the opioids into the Dubrovka Theater in an attempt to subdue Chechen resistance fighters (Riches, Read, Black, Cooper, & Timperley, 2012). Domestically, dozens of poisoning deaths among children have been credited to transdermal patches for either their mistakened application (when believed to be a band-aid), or accidental ingestion through chewing (Stoecker, Madsen, Cole, & Woolsey, 2016). The most significant aspect to this legacy by far however, has been fentanyl's role in the illicit opioid epidemic. Between 2013 and 2020, the number of opioid-related deaths attributed to synthetic opioids increased by a factor of 18, advancing to the point of accounting for 82% of all opioid-related deaths in 2020 (Hedegaard, Miniño, Spencer, & Warner, 2021). While short-comings in cause-of-death reporting keep the precise number of deaths resulting from specifically fentanyl use difficult to determine, drug seizure data from the National Forensic Laboratory Information System (NFLIS) suggest that 59% of all analgesics, and as much as 91% of non-Buprenorphine synthetic opioids seized by law enforcement contain fentanyl, an analog, or a chemical precursor used for its production (DEA, 2021). The reasons for this extraordinary change is multi-fold, but from the supply-side it's been largely driven by economic factors: Fentanyl is cheap to produce and its high potency allows for both easier cross-border movement, and cutting with other substances (Greenwood & Fashola, 2021). Accompanying fentanyl's growing prevalence within illicit drug markets has been an increased interest among toxicologists in understanding the precise hazards the opioid poses through passive exposure. In their review of the extant clinical literature, Moss et al. (2018) find little evidence to corroborate the idea that momentary contact poses

any significant health risk: The required duration of continued exposure to powdered fentanyl to achieve a therapeutic- let alone, toxic- dosage through inhalation is on the order of *hours*, not seconds. Moreover, dermal contact alone does not appear to be capable of permitting the absorption of fentanyl to the bloodstream.² In one recent noteworthy event, a first responder was exposed to a large quantity of analytically-verified liquid fentanyl when it was splashed over their skin, but exhibited no clinical effects of opioid absorption (Feldman & Weston, 2022).

Numerous harm reduction policies have been advanced in an effort to combat the worsening opioid epidemic, but few have received as much attention as increasing the availability of naloxone. As an opioid antagonist, naloxone is capable of reversing respiratory depression from acute opioid intoxication within minutes of administration. Because of its life-saving capacity, it has been recognized as an ‘essential medicine’ by the World Health Organization. Since auto-injector and intranasal naloxone devices received medical approval for emergency use in 2014 and 2015, respectively, their use has expanded significantly and are now widely issued to emergency medical services, law enforcement, fire departments, and community health clinics.

Despite naloxone’s demonstrated life-saving capabilities, questions remain on its broader efficacy in reducing opioid mortality. Empirical efforts at understanding the influence of increased accessibility of naloxone as a policy response to the opioid epidemic have mostly focused on changes to Naloxone Accessibility Laws (NAL). Rees et al. (2019) look at changes to both NALs and Good Samaritan Laws (GSL) and find that NAL adoption yields significantly negative effects on opioid mortality, but that these estimates are almost entirely driven by early-adopters, suggesting these treatments were probably endogenous responses. Conversely, Doleac and Mukherjee (2022) employ a similar research

²Moss et al. (2018) provide the following scenario to illustrate just how unlikely immediate reaction is:

If bilateral palmar surfaces were covered with fentanyl patches, it would take ~14 min to receive 100mcg of fentanyl . . . This extreme example illustrates that even a high dose of fentanyl prepared for transdermal administration cannot rapidly deliver a high dose.

That is, even when the entire surfaces of both palms are covered with patches, it still takes more than 10 minutes to receive a therapeutic dose. They note that these figures are unrealistic however, as they are “based on fentanyl patch data, which overestimates the potential exposure from drug in tablet or powder form in several ways.”

design and find evidence that naloxone access had no significant effect on opioid mortality. Erfanian et al. (2019) attempt to account for spillover effects across borders in regards to both opioid mortality and NALs by estimating a spatial Durbin model. They find NALs have very mixed results, but generally do not appear to significantly decrease mortality directly (though certain NALs yield positive and negative effects when examined in the aggregate with spillovers to neighbors). This lack of consistent or clear estimates for these potential effects highlights a common theme: Meta-analysis (Smart, Pardo, & Davis, 2021) of multiple literatures find that NALs have mixed, if only slightly-positive impacts on opioid mortality. The ambiguity here is often attributed, like other similar harm reduction policies (Packham, 2019), to offsetting moral hazard behavior (Doleac & Mukherjee, 2022), but my results here hint to the possibility of another attenuating factor. Simply making naloxone more available may be insufficient as a lifesaving measure if people are reluctant to use it.

To illustrate the role that naloxone plays in overdose situations, I describe a typical scenario and the standard protocols employed by those rendering aid. Firstly, note that the first responders to an overdose scene are often not emergency medical services (EMS), but law enforcement officers (LEO). Officer surveys and analysis of bodycam footage suggest that in the majority of cases LEO are first to the scene (Smiley-McDonald, Attaway, Richardson, Davidson, & Kral, 2022; White, Watts, Orosco, Perrone, & Malm, 2022), sometimes beating EMS by several minutes. While there is some heterogeneity with this tendency in regards to urbanicity (Smiley-McDonald et al., 2022) (officers in rural regions report arriving around the same time as EMS), even in areas where not commonly first to the scene, LEOs are still more likely to administer naloxone than other responders when they are first (Macmadu et al., 2022). Because opioid-induced respiratory failure can cause death by brain hypoxia within a matter of minutes and responders are already operating off a time delay when arriving to a scene, LEOs often immediately administer naloxone then attempt CPR, so as to “buy time” before EMS arrives (Smiley-McDonald et al., 2022). Depending on victim response, first responders may administer multiple doses of naloxone; and if stabilized, they are typically either arrested, escorted to a hospital, or

released at the scene.

Because of the likelihood of being in close proximity to narcotics, it has been recommended that first responders to suspected fentanyl overdoses don nitrile gloves and- when believed to be airborne- facemasks (Moss et al., 2018). While these recommendations do not differ materially from those made for any other drug overdose,³ it has not stopped private industry from marketing specialty fentanyl personal protective equipment (PPE). These fentanyl-proof gloves, testing equipment, and hazmat suits have been adopted by some police departments (Herman et al., 2020), but have also been panned by toxicology experts as unnecessary (Lynch, Suyama, & Guyette, 2018). Considering that the margin of time needed for an overdose to become lethal could be on the order of seconds, delaying needed aid to a victim to put on superfluous PPE has prompted calls to reconsider these practices and to relax even the standard recommendation for use of N95 respirators (Lynch et al., 2018; Winograd, Phillips, et al., 2020; Herman et al., 2020; Attaway et al., 2021).

1.2.2 The fentanyl misinformation panic and its origins

The principle vector through which fentanyl misinformation appears to be disseminated is media reporting on supposed exposure events, most typically those involving law enforcement officers. So prevalent have these media pieces become in communities hardest hit by the opioid epidemic that reported-on scenarios often follow a standard formula: Following an attempted drug possession arrest, an officer comes into contact with a powdered narcotic; through either the admission of the offender or just supposition, said officer comes to believe that this substance is fentanyl; after a period of several minutes the officer reports feelings of dizziness, shortness of breath, and may even faint; in an attempt to resuscitate the exposure victim, other officers or first responders may administer naloxone or escort them to a hospital to receive treatment. Affected first responders, their peers, and accompanying media portrayals may attest these reactions to acute opioid toxicity, but even if one were to disregard the extreme unlikelihood of

³The NIOSH recommendations are intended for when any illicit drugs are at an emergency medical scene and offer no additional considerations for fentanyl specifically. See: <https://www.cdc.gov/niosh/topics/fentanyl/risk.html>

passive fentanyl exposure eliciting such medical responses, there is virtually no evidence to corroborate the veracity of these claims (Lynch et al., 2018; Herman et al., 2020; White et al., 2022). Herman et al. (2020) combed through more than one thousand media reports involving supposed first responder fentanyl exposure events between 2014-2018 and could not find a single instance where either the affected parties reported a plausible poisoning scenario or laboratory testing confirmed poisoning. Instead, they find that the most commonly reported symptoms are consistent with stress-induced panic, and that these reactions are probably psychosomatic in origin.

Although their underlying accuracy is disputed, the media presence first responder fentanyl exposure events maintain is far from trivial. Estimates on the upper-bounds of cumulative facebook user-views that these media reports have received between 2015 and 2019 is approximately 70 million, while only 6.6% of these shares correspond to articles that correctly refute the incidental exposure hazards (Beletsky et al., 2020). Accordingly, surveys suggest that knowledge of these erroneous exposure hazards have permeated aggressively through the first responder community, with as many as 80% of queried law enforcement and emergency medical services members agreeing that momentary contact with fentanyl can be deadly (Persaud & Jennings, 2020; Del Pozo et al., 2021; Attaway et al., 2021; Berardi, Bucorius, Haggerty, & Krahn, 2021; Bucorius, Berardi, Haggerty, & Krahn, 2022). Of those law enforcement officers who echoed these sentiments, many note that they had learned of the phenomenon second-hand and not through formal police channels, suggesting that media coverage may be a contributing factor for first responder perceptions specifically (Attaway et al., 2021). Moreover, these beliefs also appear to translate directly to first responder behavior, with some law enforcement admitting to an unwillingness to render first aid to those they suspect of suffering from fentanyl poisoning (Berardi et al., 2021; Bucorius et al., 2022).

The valid hazardous concerns of fentanyl as an accidental poisoning agent can be traced to historical events with relative ease, yet the origins of the fentanyl exposure hazards myth is somewhat more opaque: Urban legends of malefactors clandestinely dosing unsuspecting

highway patrolmen date back to at least the 1970's⁴; and parallels have been noted to earlier, similarly specious medical panics regarding first responder exposure misinformation during the HIV/AIDS epidemic⁵ and early waves of clandestine methamphetamine lab raids the 1980's and 90's⁶ (Bucerius et al., 2022). What is known for certain is that beginning in 2016, medical toxicologists began receiving inquiries concerning the veracity claims made in the media that fentanyl could harm first responders on touch (Herman et al., 2020). In that same year, the US Drug Enforcement Administration, published a press release describing one exposure event involving law enforcement in New Jersey, the details of which were shared further by the National Police Foundation (Del Pozo et al., 2021).

Beyond the influence of these agency press releases, social and news media analysis performed by Beletsky et al. (2020) pinpoint one heavily-reported event as being the primary culprit in cementing these ideas within the public consciousness. While attempting to make an arrest in May 2017, an East Liverpool, OH police officer was exposed to what was believed to be fentanyl. Within minutes he became lightheaded and naloxone was administered several times, but his symptoms were severe enough to eventually require hospitalization. Media coverage of the event was swift, with early news reports receiving tens of thousands of facebook shares (Beletsky et al., 2020). As illustrated in Figure 1.1, average search prevalence among Google queries for terms related to fentanyl exposure increased by a factor of five immediately following this particular event. Similarly, in Figure 1.2 I compare search prevalence for all queries involving the term "fentanyl" across time within the Youngstown, OH Designated Market Area (DMA) (where East Liverpool is located) against neighboring DMAs, as well as all other DMAs in the surrounding states

⁴See "LSD Given to Police Officer" at Snopes here: <https://www.snopes.com/fact-check/jar-jar-drinks/>

⁵One ambulance service director in New York State, who had worked as an EMT during the HIV/AIDS crisis made this comparison more overt, saying (Bump, 2018):

It was a lot of hype ... We didn't understand it, we didn't know how it was transmitted, and I think we're seeing the same thing here. But the reality is, the initial scares about exposure to this drug just have not panned out.

⁶In a similar fashion to specialized fentanyl PPE, beliefs about long term complications related to meth lab exposure among retired law enforcement officers prompted the Utah state government to finance a controversial therapy regimen in 2007. The sauna-based therapy (which was developed by L. Ron Hubbard and delivered through a Church of Scientology-associated organization (Scientology Critical Information Directory, 2009)) claimed to "sweat out" toxins, though this was criticized for having little to no scientific basis (Bonisteel, 2015).

of Pennsylvania, Ohio, and West Virginia⁷. One can observe that in the media market local to East Liverpool, general interest in fentanyl peaks immediately following the 2017 event, before returning to similar search frequencies of the neighboring areas. This is consistent with Beletsky et al. (2020)'s observation that media coverage of this event- despite its relative popularity- varied substantially across space. Even as more events such as these would unfold involving LEOs from all over the country, the East Liverpool event remains unique in terms of both its timing, and ultimate breadth of coverage.

In 2018, the American College of Medical Toxicology (ACMT) and American Academy of Clinical Toxicology (AACT) released a joint statement to counter the sensationalist claims which had been made in the preceding two years (Moss et al., 2018), but this effort appears to have largely fallen on deaf ears. Since then, hundreds more articles have been published detailing claimed exposure events (Beletsky et al., 2020), with the phenomenon being reported as recently as January 2024.⁸

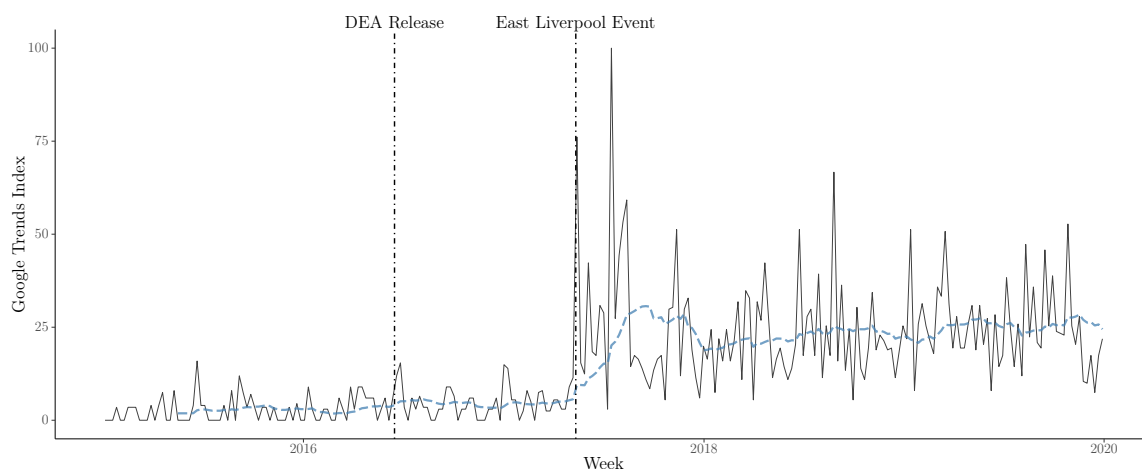
1.3 Empirical Approach

1.3.1 Data


To observe the effects of fentanyl misinformation on opioid mortality, I utilize Nielsen's Designated Media Markets (DMA) to identify treated counties within the same media market as East Liverpool, OH. As argued above, this May, 2017 event appears to be the pivotal event in establishing public misconceptions on the hazards of passive fentanyl exposure, so distinguishing between regions based on their exposure to misinformation through the media would permit one to identify treatment effects. DMAs are very similar to the Federal Communications Commission's (FCC) Television Market Areas (TMA), the legally-defined borders that determine broadcast rights and channel availability for all over-air, satellite, and cable television. TMAs borders are usually larger than actual media coverage areas (particularly in mountainous regions like Appalachia) however, so

⁷More detailed information and background on these DMAs is provided in the following section.

⁸See news article, *Florida sheriff: Deputy exposed to fentanyl, saved by Narcan*, here: <https://www.wmur.com/article/somersworth-police-fentanyl-heroin-013024/46587496>

Figure 1.1: Time series of Google search interest in the hazards of fentanyl exposure.

Google search data was collected from Google Trends and represents the relative popularity of search terms over the specified time frame; the time series was derived from querying the Google trends for “fentanyl AND (touch* OR contact* OR absor* OR inhal* OR expos*)”. Solid black line is the weekly average for the Google Trends Index, while the dashed blue line is the rolling average of the 20 preceding weeks. The dashed vertical lines demonstrate the dates of the DEA press release and the East Liverpool event.

Figure 1.2: Regional time series of Google search interest in fentanyl.

Google search data was collected from Google Trends and represents the relative popularity of search terms over the specified time frame. Time series were derived from querying the Google trends for “fentanyl” for each Designated Market Area (DMA)-quarter over 2014-2019. The Youngstown DMA contains East Liverpool, OH; while neighbor DMA includes all counties in media markets that share a border with the Youngstown DMA; and non-neighbor DMAs include all other counties in Ohio, Pennsylvania and West Virginia for which search trends data is recorded. Shaded regions represent the 95% confidence intervals and the dashed vertical line demonstrates the date of the East Liverpool event.

Nielsen DMAs are adjusted to only include counties with significant metered-viewership. I hand-code these DMAs based on a publicly available map⁹. Because there still remains the potential for some bordering counties to be treated by broadcast misinformation, I additionally employ FCC significant viewership data to include any counties that could be plausibly subject to spillovers¹⁰.

To describe mortality effects, I utilize the CDC's restricted-use multiple cause of death file. These data include the entire universe of deaths within the United States over the period 2014-2019 and are recorded at the individual-level. Following the procedures outlined by Svetla et al. (2015), I identify all overdose deaths as those with ICD-10 underlying cause of death codes X40-X44, X60-X64, X85, Y10-Y14; then identify opioid overdose deaths from these as those with with mortality-associated conditions codes T40.0, T40.1, T40.2, T40.3, T40.4, or T40.6¹¹ I aggregate these opioid overdoses and compute the per 100,000 population death rate by the county-quarter and initially drop any counties which did not record a single opioid overdose over the six year sample. To account for potential undercounting of opioid overdoses, I also use aggregated counts of overdose deaths which include the mortality-associated conditions code T50.9 for poisoning by unspecified drugs, medicaments and biological substances (Buchanich, Balmert, Williams, & Burke, 2018). As an additional robustness check, I perform placebo tests employing similar mortality rates for motor vehicle accidents, heart attacks, and assault excluding the use of drugs or medicants¹²

Additional control covariates include county-quarter demographic and economic measures, as well as opioid-use proxies such as the annual opioid dispensing rate and heroin arrest

⁹Available here: https://web.archive.org/web/20230315182138/https://thevab.com/storage/app/media/Toolkit/DMA_

¹⁰The FCC is legally obligated to conduct periodic viewership surveys to determine which specific channels receive significant viewership outside of their designated TMAs. I use the 2017 survey, which is available from here: <https://transition.fcc.gov/mb/significantviewedstations061817.pdf>

¹¹These T-codes correspond, respectively, to: opium, heroin, natural and semisynthetic opioids, methadone, synthetic opioids, and other or unspecified opioids.

¹²For motor vehicle accident deaths, I use all underlying cause of death codes corresponding to unintentional motor vehicle deaths: [V02-V04](.1-.9), V09.2, [V12-V14](.3-.9), V19(.4-.6), [V20-V28](.3-.9), [V29-V79](.4-.9), V80(.3-.5), V81.1, V82.1, [V83-V86](.0-.3), V87(.0-.8), V89.2. For assault-related deaths, I use all codes contained under X86-99 and Y00-Y05, which includes all forms of assault, excluding assault by drugs, medicaments and biological substances. For heart attack-related deaths I use all codes contained under I10-15, hypertensive diseases; I20-25, ischemic heart diseases; I46, cardiac arrest; and I50, heart failure.

rates, and policy indicators for naloxone access, good samaritan, and opioid prescription control laws. One likely confounder which may prevent the interpretation of any naive regression results as causal is the significant county-level variation in fentanyl and heroin prevalence over the observed time frame. Broader historical analyses of the opioid epidemic have emphasized two consumption innovations– the transition from prescription painkillers to heroin, and eventually from heroin to fentanyl– as epochal in defining the associated mortality (Alpert, Powell, & Pacula, 2018; Alpert, Evans, Lieber, & Powell, 2022). There is particular concern over the market transition towards primarily-synthetic opioid consumption: Because the introduction of fentanyl was so swift, disparate, and ultimately lethal, the potential for spuriously conflating those associated outcomes with media misinformation treatment effects seems a valid concern. To control for these confounding effects then, I also utilize law enforcement drug seizure data collected from Harm Reduction Ohio, which includes the entire universe of Ohio’s Bureau of Criminal Investigation’s (BCI) crime lab results for the years 2014-2019. A complete breakdown of controls employed, their sources, and spatio-temporal coverage is available in Table 1.1.

Unlike user surveys, which are largely dependent on the word of dealers in establishing the provenance and composition of traded goods, the precise chemical makeup of seized samples are determined through a gas chromatography process, and as such much less likely to omit or mistake the presence of specific opioids. The BCI laboratory is also by far the largest crime lab within the state– of Ohio’s 88 counties, only two are absent for the years observed. These data contain individual offense-level observations (including seizure data, arresting authority, and county of seizure location) and the corresponding chemical makeup of any drugs seized, which represents a significant improvement over opioid possession or intent-to-distribute arrest records. I extract from this dataset the total county-quarter counts of seizures which tested positive for fentanyl, or the closely related carfentanil, and do the same for heroin¹³.

Summary statistics for key variables of interest and covariates are listed in Table 1.2. To

¹³Because fentanyl is almost universally used as an adulterant of heroin or psychostimulants, it is worth noting that these fentanyl and heroin counts are not exclusive.

Table 1.1: Data description, coverage and sources

Data employed	Level of measure	Geographic Coverage	Temporal Coverage	Data Source
Multiple cause-of-death file	Individual deaths, aggregated to county-quarter	All US counties	2014-2019	National Vital Statistics System, Centers for Disease Control and Prevention
Percentage of the population with a credit score below 660	County-quarter	All US counties	2014-2019	Equifax Subprime Credit Population, Equifax and Federal Reserve Bank of New York
Arrests per 100k for possession or distribution of heroin and similar drugs	Month-agency counts, aggregated to county-quarter	All US counties	2014-2019	Uniform Crime Reporting, Summary Reporting System, Federal Bureau of Investigation
Unemployment rate	County-quarter	All US counties	2014-2019	Local Area Unemployment Statistics, Bureau of Labor Statistics
Percent of laborforce employed in construction	County-quarter	All US counties	2014-2019	Local Area Unemployment Statistics, Bureau of Labor Statistics
County-level demographic estimates (percent hispanic, black)	Year-quarter	All US counties	2014-2019	County Population Totals, U.S. Census Bureau
Poverty rate	Year-quarter	All US counties	2014-2019	Small Area Income and Poverty Estimates, U.S. Census Bureau
Policy indicator for whether state has a naloxone access law	State-quarter	All US counties	2014-2019	Prescription Drug Abuse Policy System
Policy indicator for whether state has a law restricting prescriptions for opioid analgesics	State-quarter	All US counties	2014-2019	Prescription Drug Abuse Policy System
Policy indicator for whether state has a drug overdose Good Samaritan Law	State-quarter	All US counties	2014-2019	Prescription Drug Abuse Policy System
Policy indicator for whether state requires the PDMP to be queried under any circumstance	State-quarter	All US counties	2014-2019	Prescription Drug Abuse Policy System
Opioid dispensing rate per 100 people	County-year	Most US counties (n=2975)	2014-2019	Centers for Disease Control and Prevention, National Center for Injury Prevention and Control
Crime lab analysis of seized drugs (percent of seizures containing heroin or fentanyl)	Individual seizures, aggregated to county-quarter	86 Ohio counties	2014-2019	Ohio Bureau of Criminal Investigation

Table 1.2: County-quarter summary statistics table.

Statistic	Control Counties		Treated Counties	
	N	Mean	N	Mean
Overdose Count	288	5.951	96	11.406
Overdoses per 100k Pop.	288	4.740	96	5.910
Annual Population	288	119,087	96	161,827
% of Pop. Hispanic	288	1.714	96	2.724
% of Pop. Black	288	4.461	96	9.438
Unemployment Rate	288	5.624	96	6.072
% of Laborforce Employed in Construction	288	4.445	96	3.734
Poverty Rate	288	13.221	96	16.058
% of Pop. with Subprime Credit	288	24.656	96	26.918
Prescription Opioid Dispensing Rate	288	82.011	96	96.646
Heroin or Related Drug Arrests per 100k Pop.	96	15.144	96	14.662
Naloxone Access Laws	288	0.931	96	0.969
Good Samaritan Laws	288	0.722	96	0.656
Opioid Prescription Restriction Laws	288	0.441	96	0.438
Mandatory PDMP Laws	288	0.913	96	0.948
Fentanyl % of Seizures	144	12.234	72	14.469
Heroin % of Seizures	144	20.507	72	24.228


All data described above represents average observed values by county-quarter for the years 2014-2019. Treated counties are defined as those within the Youngstown, OH DMA (Columbiana, Mahoning, Trumbull, OH; Mercer, PA); control counties are those that are directly adjacent and share a common border. Fentanyl and heroin seizure figures are available within Ohio alone.

align with my primary identification approach, I separate between columns 1 and 2 the statistics corresponding to the treated counties (those within the Youngstown DMA), and control counties (those within a DMA that shares a border with the Youngstown DMA). One can observe that average mortality within the treated counties is greater than their neighbors. To expand on this, in Figure 1.3 panel *a* I plot a time series of average county-month opioid overdoses for both the treated counties and their direct neighbors (that is, only the *counties* that share a border with the treated), and all other counties in Ohio, Pennsylvania, and West Virginia. Because of the relative ruralness of these treated counties, direct neighbors are probably more appropriate baselines for comparison here. The parallel trends in mortality prior to treatment, and divergence afterwards appear to lend credence to the media exposure hypothesis, although the general stabilization or decline in mortality is somewhat unexpected. In panel *b* I plot the time series for the county-month average percentages of drug seizures containing heroin and fentanyl for the entire state of Ohio. These plots demonstrate the importance of including opioid type prevalence measures, as fentanyl overtook heroin in ubiquity at almost precisely the same time as the East Liverpool event. Omission of such controls could spuriously inflate derived estimates if fentanyl was more lethal, and differentially distributed among treated and control areas.


1.3.2 Identification Strategy

My primary identification approach is to look at opioid mortality within the Youngstown Ohio DMA and compare this against the opioid mortality in some combination of bordering counties before and after the 2017 East Liverpool event. Because East Liverpool lies squarely within the Youngstown media market, and because both the timing and location of the event are seemingly random, the identifying assumption is that any difference in overdose mortality trends between counties within and outside this market can be attributed to behavioral changes among first responders due to the difference in exposure to the corresponding media coverage. In examining this particular exogenous media shock, I exploit random changes in public perceptions of the fentanyl exposure hazards to identify

Figure 1.3: Time series for opioid mortality and prevalence.



((a)) Opioid-related deaths per 100 thousand population by county-quarter. Shaded regions represent the 95% confidence intervals and the dashed vertical line demonstrates the date of the East Liverpool event.



((b)) Opioid prevalence as percent of drug seizures that tested positive for fentanyl and heroin based on the BCI data, by county-quarter. Shaded regions represent the 95% confidence intervals and the dashed vertical line demonstrates the date of the East Liverpool event.

the causal effects on first responder behavior. The counterfactual here would be that, absent some media intervention, individuals would not erroneously believe fentanyl to be so hazardous as to stymie or delay emergency response to overdoses. This is essentially the canonical differences-in-differences (DiD) research design, but I use additional data-driven methods to create a better match on pre-treatment observables,

To describe this procedure in more detail, consider a conventional approach at deriving these misinformation treatment effects using the following two-way fixed effects (TWFE) model:

$$y_{ct} = \mu + \alpha_c + \beta_t + \tau \cdot \text{MI}_{ct} \quad (1.1)$$

where y_{ct} is the opioid-related mortality rate in county c in quarter t ; α_c and β_t are county and quarter fixed effects, respectively; and MI_{ct} is an indicator equal to one representing misinformation exposure for all counties within the Youngstown DMA following the May 2017 treatment, and zero otherwise. Because this model specification holds the composition of intervention and comparison groups stable, and assuming that treatment assignment is not itself endogenous with opioid-mortality, then τ here can be interpreted as a causal average treatment effect on the treated as long as the parallel trends assumption is satisfied. A common critique with this standard DiD approach however, is that the validity of the parallel trends assumption cannot be formally tested, meaning that *ad hoc* control definitions could be yielding spurious results. Considering the relatively small size of the treated sample examined here, this is particularly threatening to the causal interpretation as the control group could be misspecified.

With this in mind, I instead estimate a synthetic differences-in-differences (SDiD) model that minimizes the following error:

$$\arg \min_{\alpha, \beta, \mu, \tau} \left\{ \sum_{c=1}^C \sum_{t=1}^T (y_{ct} - \mu - \alpha_c - \beta_t - \tau \cdot \text{MI}_{ct})^2 \hat{\omega}_c^{\text{sdid}} \hat{\lambda}_t^{\text{sdid}} \right\} \quad (1.2)$$

where $\hat{\omega}_c^{\text{sdid}}$ is a vector of statistically-derived county control weights, and $\hat{\lambda}_t^{\text{sdid}}$ is a vector of time weights computed according to Arkhangelsky et al. (2021). The SDiD control

generated by these weights minimizes the error in pre-treatment trends when compared with the treated, so that it represents a more realistic counterfactual than any *ad hoc* specification. Arkhangelsky et al. (2021) describe this as a “generalized” differences-in-differences model, because- unlike the canonical DiD which assigns uniform weight to each pre-period control- SDiD weights those control observations which best construct parallel pre-trends. Because of this strength, SDiD method has been demonstrated to generally outperform both conventional TWFE and SDiD control estimation approaches (Arkhangelsky et al., 2021). In addition to these approaches, I estimate a staggered-SDiD that considers all counties within DMAs that had at least one news article published making reference to the East Liverpool event as treated, and their bordering DMAs as controls. This is a similar identification approach to the primary method, but permits DMAs other than just Youngstown to be treated and uses variation in timing of media coverage.

A growing concern among applied researchers with interpreting either conventional synthetic control or SDiD estimates derived from observations of the dependent variable alone as causal is that the asymptotic irrelevance of auxiliary covariates may not necessarily hold over finite sample spaces. While the recommendation is often overlooked, Abadie et al. (2010) and others (Kaul, Klößner, Pfeifer, & Schieler, 2022) propose the inclusion of potential confounders whenever researchers find them relevant. Recent simulation results indicate that omitting these confounders in model specifications could not only introduce significant bias to estimates but also render the precise direction of this bias virtually unknowable from the outset (Pickett, Hill, & Cowan, 2022). As such, when possible I duplicate all estimation procedures with and without the inclusion of plausible confounders, and display both sets of results.

To utilize contemporaneous observations of relevant covariates within the synthetic differences-in-differences estimation, Arkhangelsky et al. (2021) propose first regressing the dependent variable on the covariates, then running the SDiD procedure on the obtained residuals. Kranz (2022) demonstrate however, that in instances where covariates have time-varying influence on the dependent variable, this residuals approach often fails at

constructing a SDiD control which satisfies the parallel trends condition. They instead suggest a correction approach that utilizes fitted values for the dependent variable derived from a two-way fixed effects regression including the covariates. Because several important control variables within my model- in particular, those pertaining to opioid demand- likely differ in their influence on opioid-related mortality across time,¹⁴ I opt to utilize this ‘projected’ covariate approach when running my regressions.

Lastly, it is important to acknowledge that- while initially untreated- the control counties are susceptible to contamination over time as information spillovers occur through social and national media coverage of later, similar misinformation-triggering events. It is unlikely that my model is capable of entirely controlling for spillovers. Instead, I opt to place particular emphasis on examining dynamic trends in mortality for the bulk of my analysis: If delineating between treated and untreated counties becomes more difficult as time passes, I should still be able to observe any eventual convergence between the treated and controls.

1.4 Results

1.4.1 Primary Results

The primary results of my regressions are outlined in Table 1.3. As detailed already, I estimate SDiD models which assume all counties within the Youngstown DMA are treated through exposure to misinformation following the coverage of the East Liverpool event, while donors to the SDiD control are drawn from the immediately bordering counties to this treated DMA. In the upper panel, I estimate SDiD models both with and without auxiliary covariates using several measurements for accidental drug poisoning. These estimates comprise of all drug-related poisonings; opioid-related poisonings; possibly-opioid-related poisonings which include all opioid-related deaths, as well as those coded as

¹⁴For instance, I employ possession and distribution arrests data as a proxy for illicit opioid demand, but these data do not distinguish between opioid types. Since the observed timeframe also coincides with the transition from predominantly heroin to fentanyl use, parametrizing this demand effect as constant while there are unobserved changes in illicit opioid lethality could bias pre-trends.

related to unspecified drugs; illicit opioid-related poisonings (those associated with either heroin or a synthetic opioid, like fentanyl); and synthetic opioid-related poisonings. In the lower two panels I re-estimate my opioid-related SDiD model, but stratify by decedent demographics and age groups.

Summarizing these results, I observe large and statistically-significant increases to overdose mortality rates within the Youngstown DMA following the East Liverpool event. Across specifications both with and without the inclusion of control variables, I observe overall accidental drug poisoning and opioid-related poisonings are increasing by between 2-4 additional deaths per 100,000 population. My preferred outcome measure for this analysis, opioid-related mortality, increased by approximately 2.84 deaths per 100,000 population. This jump in mortality represents a 56.6% increase to the average quarterly mortality for the treated counties relative to pre-treatment period. The results for possibly-opioid related, illicit opioid-related, and synthetic opioid-related overdose deaths highlight data quality concerns that have been voiced by other researchers. It appears likely that-consistent with earlier findings (Buchanich et al., 2018)- many illicit opioid-related overdose deaths are being coded under this general “unspecified drugs” category. A lack of adequate toxicology screening, particularly early on when fentanyl first entered illicit drug markets, may have failed to correctly identify synthetic opioids when they were in fact present. While I cannot be certain that all of these deaths represent fentanyl, or even opioid-related deaths, concern for categorical undercounting here seems valid. Because of this, caution is warranted in interpreting these more drug-specific mortality coefficient estimates, as they are almost certainly biased downward. More reassuringly however, the stratified SDiD results illustrate sensible heterogeneity in changes to opioid-related mortality. Consistent with other research on the opioid epidemic, the bulk of these effects are being driven by white males between the ages of 25-34.

To better illustrate the dynamic trends of this phenomenon, I plot time series of the observed treated opioid-related mortality against the computed counterfactuals for the estimated opioid-related mortality SDiD model in Figure 1.5. Interestingly, opioid mortality appears to decline immediately following treatment for all counties, but critically, the

Table 1.3: SDiD coefficient estimates for mortality by drug type, stratified by demographics and age.


	Pre-treatment Mean	SDiD estimates	SDiD estimates w/ covariates
Opioid-related	5.012	2.353*** (0.811)	2.838*** (1.099)
Drug-related	7.890	3.347*** (0.914)	3.953*** (1.073)
Possibly opioid-related	7.251	3.129*** (0.823)	4.051*** (1.091)
Illicit opioid-related	3.811	1.520 (1.009)	1.876 (1.201)
Synthetic opioid-related	2.221	2.322* (1.222)	2.115* (1.200)
Opioid related, Male	5.224	1.858* (1.109)	3.285** (1.552)
Opioid related, Female	5.200	1.456 (1.135)	1.720 (1.356)
Opioid related, White	3.213	2.040** (0.903)	2.849** (1.262)
Opioid related, Black	7.018	-4.034 (4.178)	-1.447 (5.876)
Opioid related, Hispanic	3.090	4.324 (2.639)	1.121 (4.016)
Opioid related, Age <25	1.398	-0.515 (1.024)	0.172 (1.042)
Opioid related, Age 25-34	13.395	6.659 (4.945)	11.445** (4.556)
Opioid related, Age 35-44	11.095	7.970* (4.611)	6.284 (4.438)
Opioid related, Age 45-54	8.129	-0.425 (1.881)	0.953 (2.052)
Opioid related, Age >54	2.258	0.767 (0.811)	0.800 (0.878)
Observations		384	384

Note:

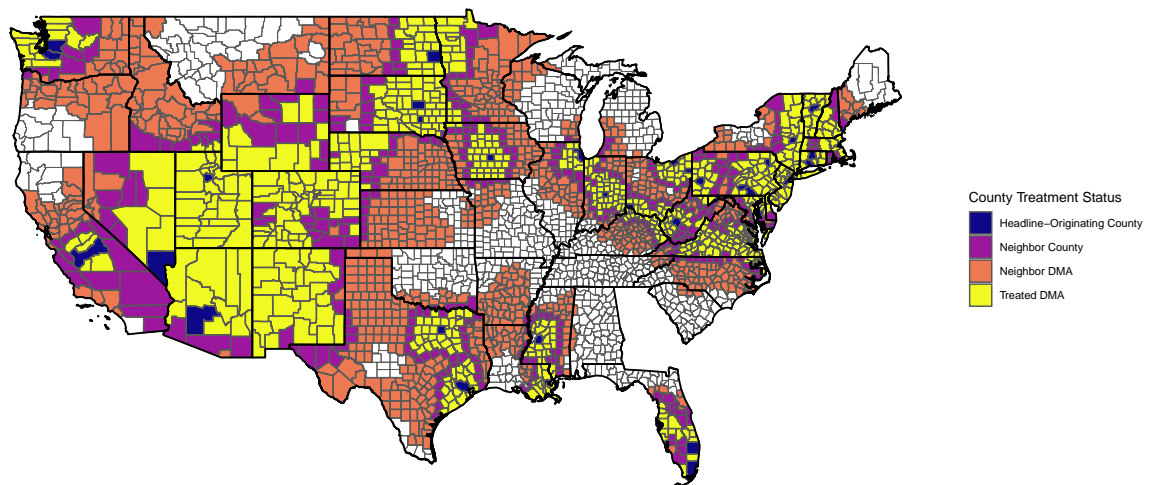
*p<0.1; **p<0.05; ***p<0.01

Results depicted here are derived by estimating an SDiD model which considers counties within the Youngstown OH DMA after the 2017 East Liverpool event as treated, and their immediately bordering counties as the SDiD control donor set. Dependent variables include per 100 thousand mortality rates for: all drug-related poisonings; opioid-related poisonings; possibly-opioid-related poisonings which include all opioid-related deaths, as well as those coded as related to unspecified drugs; illicit opioid-related poisonings (those associated with either heroin or a synthetic opioid, like fentanyl); and synthetic opioid-related poisonings. Results in column 3 are estimated by employing the time-variant covariate correction from Kranz (2022) and include the following auxiliary covariates: Percent of county population hispanic, black, or with a subprime credit score; percent of county laborforce employed in construction; unemployment and poverty rates; annual prescription opioid dispensing rate; arrests for heroin or related drugs per 100 thousand population; and policy indicators for whether the county's state had enacted naloxone access, good samaritan, mandatory PDMP, or opioid prescription restriction laws. Figures enclosed in parentheses are cluster bootstrap standard errors.


Figure 1.4: Treatment maps for East Liverpool event.



((a)) A map of counties considered as treated under my primary identification strategy. Blue-shaded region is Columbiana county, where East Liverpool is located; yellow-shaded regions are other counties within the same DMA; purple-shaded counties are bordering counties within neighbor DMAs; and the orange-shaded region are all other counties within the neighboring DMAs.



((b)) A map of national media coverage of East Liverpool event, as based on the data from Beletsky et al. (2020). Blue-shaded regions correspond to counties where a news article which made reference to the East Liverpool event originated; yellow-shaded regions are other counties within the same DMAs as those which originated coverage of the East Liverpool event; purple-shaded counties are bordering counties within neighbor DMAs; and the orange-shaded region are all other counties within the neighboring DMAs.

Figure 1.5: Plotted primary SDiD results.

((a)) SDiD results depicted here are according to the opioid-related mortality specifications from table 3. Pink-shaded regions at the bottom of plots depict time-weights, black arrowed-line represents the average treatment effect. The bottom panel is estimated employing the time-variant covariate correction from Kranz (2022) (hence the seemingly-negative values for the SDiD control the beginning of the observation period) and includes the following auxiliary covariates: Percent of county population hispanic, black, or with a subprime credit score; percent of county laborforce employed in construction; unemployment and poverty rates; annual prescription opioid dispensing rate; arrests for heroin or related drugs per 100 thousand population; and policy indicators for whether the county’s state had enacted naloxone access, good samaritan, mandatory PDMP, or opioid prescription restriction laws.

treated counties rebound much more quickly. This could be representing an incubation period for the misinformation to disseminate and take hold.

One potential confounder missing from these model specifications is the variation in the ubiquity of particularly potent illicit opioids, fentanyl and heroin, over the observed period. Prior research examining Ohio over this same timeframe (Peterson et al., 2016; Zibbell et al., 2022) has identified the regional prevalence of fentanyl as a significant driver of opioid mortality, so the inclusion of some measure of this within the estimated model appears justified. By employing drug seizure data from Ohio’s Bureau of Criminal Investigation’s crime lab (BCI), I use county-quarter counts of total seizures that tested positive for these compounds to proxy for their prevalence within local drug markets. A problematic factor with using these direct seizure counts is that they could conceivably be endogenous with

opioid deaths.¹⁵ As an alternative then- and since I am only interested in the *relative* prevalence of these opioids- I divide these drug-specific counts by the total count of all drug seizures conducted within that county-quarter. To test whether these proportional estimates are endogenous, I estimate a series of simple two-way fixed-effects models where I regress total drug seizures, fentanyl as percent of seizures, and heroin as percent of seizures on the one-year lag of opioid deaths. The results for these models are listed in the upper panel of Table 1.4, but to summarize: As anticipated, deaths do appear to be significantly decreasing the number of seizures performed in the subsequent years, while the relative proportions of these seizures being either fentanyl or heroin do not seem to be affected.

With these prior results in mind, I attempt to control for opioid prevalence variation by including the proportional measures of fentanyl and heroin ubiquity as additional covariates and re-estimate my primary SDiD model specification on the subsample of Ohio counties. While not a direct threat to the validity of my reduced form estimates, a concern with interpreting these results could be a misidentification of the underlying mechanisms. For instance, it may be the case that these misinformation shocks are increasing opioid-related mortality, but are doing so by increasing consumer demand for fentanyl by users. To test this, I replicate my Ohio SDiD results twice more, but with fentanyl and heroin prevalence on the left-hand side.

The results for the three Ohio models are listed in the lower panel of Table 1.4. In column 1, I note that even with the inclusion of fentanyl and heroin drug seizure proportions, the primary specification SDiD results within Ohio do not significantly change. In columns 2 and 3, I see that when treating fentanyl and heroin prevalence as the dependent variable, there is no significant change following the East Liverpool misinformation shock. This highlights that changes to opioid overdose death are likely not arising from an increase in demand following media reports on fentanyl's potency. Taken together, these results strengthen the central argument that these observed changes to opioid-related mortality

¹⁵This could arise, for instance, when a year of unexpectedly high opioid-related deaths within a county prompts local policymakers to invest more heavily in drug enforcement, and consequently sees an increase in seizures performed in the subsequent years.

Table 1.4: Ohio fentanyl and heroin prevalence analysis results.

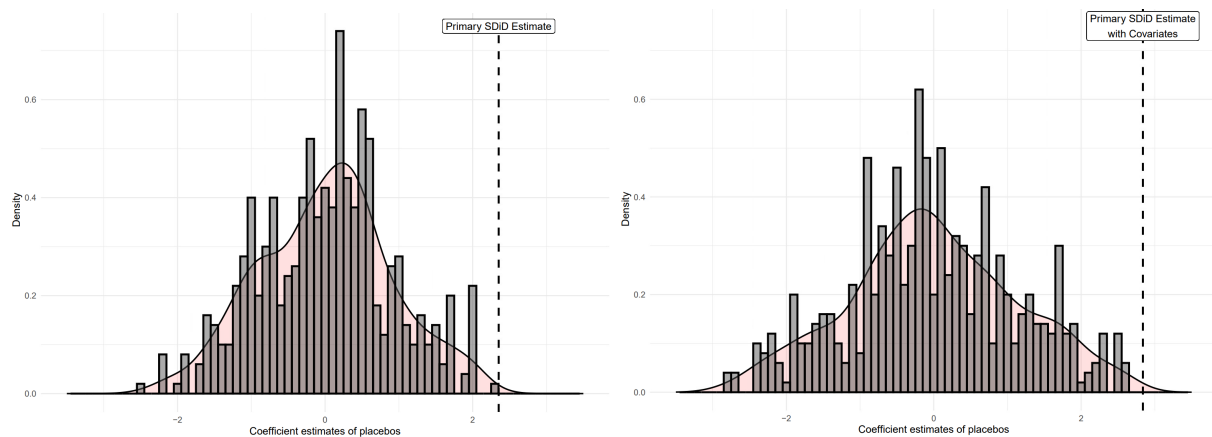
<i>(a) Opioid prevalence endogeneity test</i>			
	<i>Dependent variable:</i>		
	Total Drug Seizures	Fentanyl Seizures, % of Total	Heroin Seizures, % of Total
	(1)	(2)	(3)
Lagged Opioid Overdose Deaths	-0.290** (0.124)	0.001* (0.0003)	0.00000 (0.0004)
Observations	1,740	1,740	1,740
<i>(b) SDiD results with opioid prevalence measures</i>			
	<i>Dependent variable:</i>		
	Opioid-Related Overdose Deaths Per 100k	Percent of Drug Seizures Containing Fentanyl	Percent of Drug Seizures Containing Heroin
SDiD Estimates	3.190** (1.315)	0.015 (0.036)	0.101 (0.073)
Pre-treatment Mean	5.821	0.075	0.311
Observations	216	216	216

Note:

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

((a)) The results in the top panel are derived by regressing each Ohio county-quarter measure of the dependent variables for the years 2015-2019 on the number of opioid-related deaths that occurred in the same county-quarter of the preceding year; as well as including county and quarter fixed-effects for all Ohio counties.

((b)) The results in the bottom panel outline SDiD coefficient estimates derived from the primary model specification on the Ohio subsample, along with their bootstrapped standard errors in parentheses. Column 1 replicates the primary opioid-related results estimate from table 3 including covariates, but additionally includes measures for the percent of drug seizures conducted within those county-quarters that tested positive for heroin and fentanyl. Columns 2 and 3 follow the same controls specification, but set the fentanyl and heroin drug seizure percentages as the dependent variable.

Figure 1.6: Distribution of placebo test results.

Histogram and density plots above describe the empirical noise distribution for the primary opioid-related mortality SDiD model controls. Coefficient estimates here are derived according to the placebo protocol outlined by Arkhangelsky et al. (2021), which randomly assigns controls as treated and the SDiD model is re-estimated on the donor set alone. Distributions are based on 500 replications for each model specification. Dashed lines indicate the value of the estimated SDiD treatment effects.

are being driven by some external factor other than fentanyl, or even heroin prevalence.

1.4.2 Robustness Checks

An initial concern with my estimates is that cluster bootstrap-derived standard errors are less dependable for small treated sample sizes. Because I have only four treated counties, I re-estimate my primary SDiD specifications, but instead employ the placebo protocol outlined by Arkhangelsky et al. (2021). This approach is similar to permutation tests performed in randomization inference used for conventional DiD estimators (Conley & Taber, 2011): To directly estimate the noise level of the control units, a number of controls are randomly assigned as treated and the SDiD model is re-estimated on the donor set alone. Assuming homoscedasticity across units, this variance estimator would provide more accurate- if also more conservative- confidence bounds for the causal treatment effect. I perform this placebo procedure using 500 random placebos and plot the empirical distribution for their derived SDiD coefficients in Figure 1.6. Across specifications both with and without the inclusion of covariates, I find my initial SDiD estimates for opioid-related mortality retain their 99% significance level .

I re-estimate my opioid-related mortality model with considerations for a spate of other potential threats to validity and list the results in Table 1.5. These variations include an alternative treatment specification meant to control for information spillovers that includes any counties the FCC has listed as having significant viewership of any stations within the Youngstown DMA; alternative dependent variables of opioid-related death counts in levels and logs derived from the inverse hyperbolic sine transformation, rather than mortality rates; and a conventional DiD model. In the top panel I estimate these across all observations for the period 2014-2019, while in the bottom panel I re-estimate the opioid-related mortality model on the subset of observations occurring after the October 2015 adoption of the ICD-10 coding system to account for any potential data inconsistencies. For each of these, I experiment with several different donor-set specifications to derive my SDiD controls. Under my preferred specification in column 1, donors to the SDiD control are drawn from the immediately bordering counties to the treated DMA; under the specification in column 2, I expand this donor set to include all counties of bordering DMAs; and in column 3 I include all counties in bordering DMAs but exclude immediately bordering counties so as to control for spillovers. In Table 1.6 I replicate these results, but use local commuter zone delineations from Fowler and Jensen (2020) instead of DMAs. Under the primary donor set specification, the magnitude of the coefficients listed in column 1 are consistent with the preferred SDiD estimates, and are broadly significant. DiD estimates are qualitatively similar to SDiD, but are insignificant, which highlights the potential advantages that this more generalized estimation approach affords. Coefficients generally maintain their magnitudes across the wider donor set definitions in columns 2 and 3, but are noisier.

To better understand these results, I plot the three SDiD controls based on the different donor set specifications from row 1 of Table 1.5 in Figure 1.7. One can observe that the generated SDiD controls are nearly identical across these specifications, but that the alternative donor set definitions including more counties from the bordering DMAs are weighting earlier observations from the pre-period more heavily (represented by the shaded regions in the bottom-left). It is primarily because of these differing time weights-

Table 1.5: SDiD robustness checks, DMAs as treatment unit


	<i>SDiD control donor set:</i>		
	Adjacent counties as control	Bordering DMAs as control	Bordering DMAs without adjacent counties
<i>Estimates without Covariates</i>			
Youngstown DMA as treated	2.353*** (0.811)	1.352 (1.042)	1.477 (1.168)
Significant viewership of Youngstown DMA station as treated	2.031** (0.848)	1.643* (0.986)	1.671* (0.965)
Opioid-related death rate, logged	0.417** (0.186)	0.161 (0.211)	0.222 (0.315)
Opioid-related deaths, levels	3.600** (1.668)	3.772 (2.229)	4.049* (2.136)
DiD	1.493 (0.980)	1.536 (0.987)	1.551* (0.883)
<i>Estimates with Covariates</i>			
Youngstown DMA as treated	2.838*** (1.065)	0.891 (0.801)	0.818 (0.837)
Significant viewership of Youngstown DMA station as treated	2.394* (1.334)	1.215* (0.717)	1.088 (0.765)
Opioid-related death rate, logged	0.519** (0.233)	0.040 (0.145)	0.014 (0.187)
Opioid-related deaths, levels	4.138*** (1.579)	2.935** (1.302)	3.176* (1.745)
DiD	2.214* (1.134)	1.078* (0.582)	0.920 (0.685)
Observations	384	1,200	912
<i>Estimates without Covariates</i>			
Post ICD-10 adoption	1.994** (0.926)	1.365 (0.978)	1.473 (1.055)
<i>Estimates with Covariates</i>			
Post ICD-10 adoption	3.091* (1.681)	1.429 (0.898)	1.390* (0.766)
Observations	272	850	646

Note:

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

SDiD results for opioid-related mortality depicted here include: An alternative treatment specification that counts any counties the FCC has listed as having significant viewership of any stations within the Youngstown DMA as treated; alternative dependent variables of opioid-related death counts in levels and logs derived from the inverse hyperbolic sine transformation; and a conventional DiD model. The bottom panel re-estimates the opioid-related mortality model on the subset of observations occurring after the October 2015 adoption of the ICD-10 coding system. Column 1 defines donors to the SDiD control from the immediately bordering counties to the treated DMA; column 2, expands this donor set to include all counties of bordering DMAs; and in column 3 includes all counties in bordering DMAs but excludes immediately bordering counties. The covariate-inclusive estimates employ the correction from Kranz (2022) and includes the following time-variant controls: Percent of county population hispanic, black, or with a subprime credit score; percent of county laborforce employed in construction; unemployment and poverty rates; annual prescription opioid dispensing rate; arrests for heroin or related drugs per 100 thousand population; and policy indicators for whether the county's state had enacted naloxone access, good samaritan, mandatory PDMP, or opioid prescription restriction laws.

Figure 1.7: Comparison of SDiD controls constructed under different donor set specifications.



((a)) SDiD results depicted here are according to the opioid-related mortality specifications from row 1 of table 5. Shaded regions at the bottom of plots depict time-weights for the corresponding donor set specification, which include: All immediately adjacent counties to the Youngstown OH DMA, all counties in within bordering DMAs, and all counties in within bordering DMAs excluding immediately adjacent counties to the Youngstown OH DMA, respectively.

not the composition of the SDiD control- that the alternative donor set specifications are yielding smaller, noisier treatment effect estimates. SDiD time weights are assigned by minimizing the error between pre- and post-treatment observations of the dependent variable on all controls (that is, it affords more weight to pre-treatment periods which are better predictors of post-treatment control outcomes). If the donor set contains controls which are wholly inappropriate for construction of the SDiD control, the unit weighting algorithm would assign them low weights when estimating treatment effects; but because the time weighting algorithm is applied across *all* members of the donor set, then these invalid controls are receiving the same weight as any other. Therefore, when appropriate and invalid controls follow different time trends, the vector of generated time weights is probably biased away from being the best predictor of relevant post-treatment outcomes. Put differently, the inclusion of control counties which are qualitatively different from those within the treated DMA could be excessively weighting early-period observations if these inappropriate control counties report flat or declining opioid mortality rates, rather

than the wider increasing trend. One way of possibly identifying this biasing effect is to examine covariate balance on the SDiD-weighted controls against the treated counties. The reasoning behind this is that a donor set with better balance in terms of observables linked to opioid mortality should result in a SDiD control that more accurately mirrors an ideal counterfactual. I perform a series of covariate balancing tests for each of these three donor set specifications and list the results in Table 1.7. While the weighted controls for the larger donor sets do appear to be better balanced for some county characteristics, I find that for essential measures related to opioid use—opioid overdose, dispensing, and arrest rates, as well as policy status for NALs and mandatory PDMP reporting – my primary specification demonstrates a greater balance than the alternatives. This is intuitive when considering Youngstown’s locale: As the radius of counties included within the donor set is expanded outward, it begins to encroach on the denser, more urban Cleveland, Akron and Pittsburgh metropolitan areas. Nevertheless, the magnitude of the alternative SDiD estimates do not vary substantially, so taken together with these other considerations I retain the initial, adjacent counties specification as my preference.

Because my principal identification strategy considers only one relatively small media market as treated, it is possible that these results could be driven by some unobserved change to the underlying first responder mechanism other than the misinformation effect that I describe. For instance, it could be that counties within the Youngstown DMA experience similar changes in law enforcement or EMS staffing and response policies that incidentally coincide with the East Liverpool event. However unlikely, in such an instance my estimates would be sizably biased upwards. To descriptively analyze this possibility, I collect municipal- and county-level expenditures data on police protection, fire prevention, and health services spending from the Census’ Annual Survey of State and Local Government Finances¹⁶. I plot time series for per-capita spending in Figure 1.8 and compare the expenditures made within Youngstown OH DMA against those made by governments elsewhere in Ohio, Pennsylvania, and West Virginia. I observe no substantial

¹⁶Depending on the specifics of local and county government program structure, outlays for first responder services could appear in any one of these categories. Generally, since ambulance services are largely operated as private entities in rural areas, most government expenditures for emergency medical training and equipment will appear as either police protection or fire prevention expenses

Table 1.6: SDiD robustness checks, commuting zones as treatment unit

	<i>SDiD control donor set:</i>		
	Adjacent counties as control	Bordering commuter zones as control	Bordering commuter zones without adjacent counties
<i>Estimates without Covariates</i>			
Youngstown commuter zone as treated	2.031** (0.813)	1.809** (0.893)	1.680 (1.076)
Opioid-related death rate, logged	0.290 (0.185)	0.219 (0.191)	0.249 (0.250)
Opioid-related deaths, levels	2.979* (1.557)	3.463* (1.874)	3.721* (1.945)
DiD	1.562* (0.868)	1.884** (0.752)	2.032*** (0.782)
<i>Estimates with Covariates</i>			
Youngstown commuter zone as treated	2.394*** (1.190)	1.534* (0.884)	1.187 (0.874)
Opioid-related death rate, logged	0.457* (0.269)	0.122 (0.186)	0.006 (0.235)
Opioid-related deaths, levels	3.329*** (1.547)	3.113* (1.718)	3.676 (2.250)
DiD	1.933 (1.199)	1.636** (0.717)	1.693** (0.778)
Observations	384	960	696

Note:

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

SDiD results for opioid-related mortality depicted here substitute local commuter zone delineations from Fowler and Jensen (2020) for DMAs and include the following variations: Alternative dependent variables of opioid-related death counts in levels and logs derived from the inverse hyperbolic sine transformation; and a conventional DiD model. Column 1 defines donors to the SDiD control from the immediately bordering counties to the treated commuter zone; column 2, expands this donor set to include all counties of bordering commuter zone; and in column 3 includes all counties in bordering commuter zone but excludes immediately bordering counties. The covariate-inclusive estimates employ the correction from Kranz (2022) and includes the following time-variant controls: Percent of county population hispanic, black, or with a subprime credit score; percent of county laborforce employed in construction; unemployment and poverty rates; annual prescription opioid dispensing rate; arrests for heroin or related drugs per 100 thousand population; and policy indicators for whether the county's state had enacted naloxone access, good samaritan, mandatory PDMP, or opioid prescription restriction laws.

relative trend changes within the Youngstown OH DMA following the East Liverpool event. Because of the inconsistency in spatial coverage for these data however (not all local and county governments surveyed in every year), I would emphasize that these figures are only suggestive.

To address this issue then, I estimate three additional model sets identical to my preferred specification that instead use motor vehicle accidents, heart attack-related, and assault-related death rates as the dependent variable. The rationale here is that if there is some alternative factor influencing first responder behavior, one would be able to observe similar changes in other common forms of death where mortality is subject to these agents' behavior (that is, other causes of death where lives can be saved with timely intervention by LEO and EMS). If not however, these regressions would act as falsification tests and return null results. The results of these falsification tests are outlined in Table 1.8, and as expected, motor vehicle, heart attack and deaths show no significant variation within the Youngstown DMA compared to control areas.

1.5 Mechanisms Analysis

1.5.1 Identifying media's direct role in misinformation shocks

Though there is evidence to corroborate the claim that the East Liverpool event influenced opioid mortality, it is still unclear how precisely this occurred. I have assumed up to this point that media coverage is the primary driver, but social media and word-of-mouth are equally-plausible vectors for misinformation. To test this, I employ a slightly different treatment definition which utilizes fentanyl misinformation media coverage data collected by Beletsky et al. (2020) to identify the associated effects of mass media reporting on mortality. These data collected from the Mediacloud system cover archived news articles for the period 2015-2019 which contain various combinations of phrases indicating the presence of erroneous fentanyl exposure hazards information. ¹⁷

¹⁷Beletsky et al. (2020) manually confirmed the content of each entry as misinformation and code for each observation the date of publication, location of publisher, associated event (for instance, whether the article makes explicit reference to the East Liverpool or some other first responder incident). I extract

Table 1.7: Covariate balance tables for SDiD donor set specifications.

Covariate	<i>SDiD control donor set:</i>								
	Adjacent counties as control			Bordering DMAs as control			Bordering DMAs without adjacent counties		
	Control	Treated	Adjusted	Control	Treated	Adjusted	Control	Treated	Adjusted
	Weighted Means	Weighted Means	Mean Difference	Weighted Means	Weighted Means	Mean Difference	Weighted Means	Weighted Means	Mean Difference
Overdoses per 100k Pop.	5.181	5.204	-0.050	4.499	5.012	0.185	4.356	5.007	0.158
Annual Population	119,485	162,781	0.559	152,957	163,209	0.078	164,715	163,246	-0.042
% of Pop. Black	4.391	9.385	1.196	5.301	9.362	0.706	5.626	9.361	0.566
% of Pop. Hispanic	1.651	2.613	0.689	2.033	2.581	0.278	2.175	2.579	0.159
Unemployment Rate	6.264	6.589	0.193	6.364	6.471	0.045	6.727	6.806	0.020
% of Laborforce Employed in Construction	4.086	3.502	-0.328	4.534	3.586	-0.512	4.735	3.563	-0.580
Poverty Rate	13.373	16.456	1.080	14.523	16.525	0.751	14.890	16.557	0.648
% of Pop. with Subprime Credit	25.024	27.305	0.654	25.897	27.552	0.542	26.286	27.665	0.512
Prescription Opioid Dispensing Rate	87.896	104.463	0.849	82.597	105.887	1.156	80.348	105.893	1.331
Heroin or Related Drug Arrests per 100k Pop.	15.983	15.104	-0.085	15.368	15.192	-0.108	15.790	15.216	-0.130
Naloxone Access Laws	0.938	0.969	0.001	0.858	0.942	0.112	0.756	0.893	0.116
Good Samaritan Laws	0.562	0.406	-0.148	0.490	0.365	-0.086	0.420	0.357	-0.085
Opioid Prescription Restriction Laws	0.052	0.031	-0.022	0.025	0.019	-0.009	0.042	0.036	-0.005
Mandatory PDMP Laws	0.948	0.969	0.001	0.853	0.904	0.066	0.803	0.857	0.063

Figures in the table above represent the weighted means and standardized mean differences on observables when comparing covariate values for the treated counties against the weighted controls drawn from the SDiD control donor set. Columns 1-3, represent the donor set specifications for all immediately adjacent counties to the Youngstown OH DMA; columns 4-6, all counties in within bordering DMAs; and columns 7-9, all counties in within bordering DMAs excluding immediately adjacent counties to the Youngstown OH DMA. Weights assigned to controls are derived from the opioid-related mortality SDiD model estimated without covariates.

Table 1.8: Falsification test SDiD results.

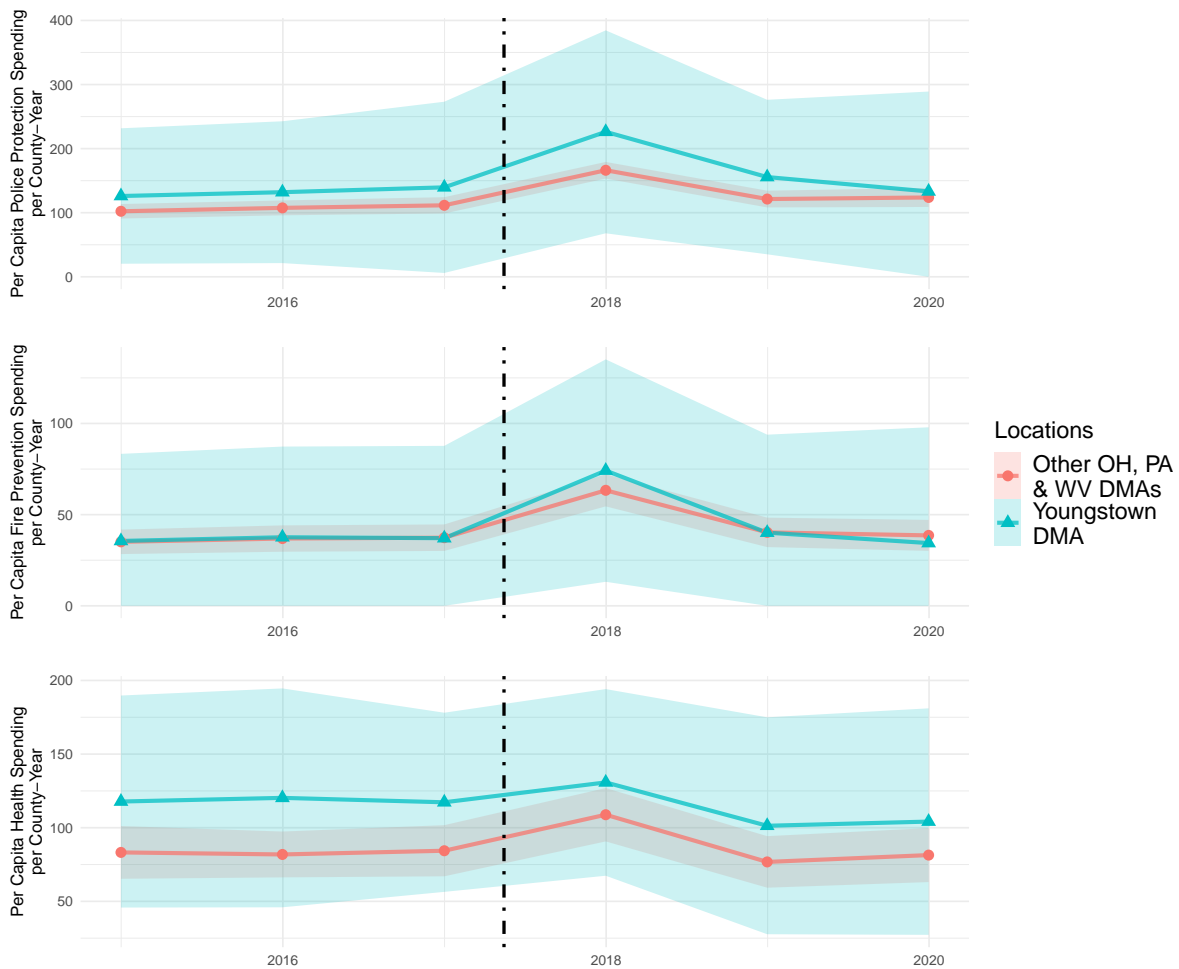
	<i>Dependent variable:</i>		
	Motor Vehicle Accident Deaths Per 100k	Heart Attack-Related Deaths Per 100k	Assault-Related Deaths Per 100k
Without Covariates	-0.484 (0.557)	-0.144 (4.695)	0.076 (0.260)
With Covariates	-0.703 (0.648)	3.938 (6.321)	0.026 (0.293)
Pre-treatment Mean	2.605	66.37	1.152
Observations	384	384	384

Note:

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

The results above are derived by following the same specifications as the opioid-related mortality models in table 3, row 1, but using mortality rates for motor vehicle accident, heath attack, and assault-related deaths. Auxiliary covariates include: Percent of county population hispanic, black, or with a subprime credit score; percent of county laborforce employed in construction; unemployment and poverty rates; annual prescription opioid dispensing rate; arrests for heroin or related drugs per 100 thousand population; and policy indicators for whether the county's state had enacted naloxone access, good samaritan, mandatory PDMP, or opioid prescription restriction laws. Figures enclosed in parentheses are cluster bootstrap standard errors.

Figure 1.8: Regional time series of local and county-level government expenditures on police protection, fire prevention, and health services.



((a)) Time series figures for per capita municipal- and county-level expenditures on police protection, fire prevention, and health services spending from the Census’ Annual Survey of State and Local Government Finances. Plots are averages across counties within the Youngstown OH DMA against all other counties in Ohio, Pennsylvania, and West Virginia by county-year. Shaded regions represent the 95% confidence intervals and the dashed vertical line demonstrates the date of the East Liverpool event.

I estimate these media-exposure models employing the staggered SDiD method outlined by Arkhangelsky et al. (2021) in their appendix. This approach separates treated groups by their treatment date, estimates an SDiD model for each treated group on the pooled control units (while excluding the other treated units), then generates the average treatment effect on the treated (ATT) as a weighted average of each sub-group's ATT according to their relative proportion of post-treatment observations. The treatment here is defined as whether a given county's DMA has originated some media coverage prior to the observed date, while controls apply the methodology of my primary estimates and are drawn from a donor set of all directly adjacent counties to treated DMAs. Because these models are being identified on regional variation in reporting alone and include time fixed effects, these estimates would correspond to only the influence of local, rather than national reporting by media outlets on opioid mortality. Following (Packham, 2019), I increase this restrictiveness when performing this national-scale analysis to include only counties which recorded at least one opioid overdose for each year in the sample. A map of the misinformation-originating counties, their DMAs and neighboring DMAs based on the Beletsky et al. data is depicted in panel (b) of Figure 1.4.

The results of my media exposure staggered-SDiD are listed in Table 1.9. In column 1, I use the complete set of all treated DMAs, while in column 2 I exclude the Youngstown OH DMA to account for potentially-biasing local misinformation vectors (e.g. word-of-mouth deriving from those involved or otherwise familiar with the East Liverpool event, absent any media coverage). Though attenuated downward relative to estimates which consider only the Youngstown OH DMA as treated, opioid-related mortality rates demonstrate significant increases within media markets following reporting on the 2017 East Liverpool event. Under the preferred staggered-SDiD specification which excludes the Youngstown OH DMA and includes auxiliary covariates, I observe an increase of 0.448 opioid-related deaths per 100 thousand people, or an increase of approximately 12%. Even when restricting the treatment definition to regional media reporting, these specific misinformation shocks are consistently increasing opioid-related mortality at a substantial

from these data all coverage relating to the East Liverpool event nationally and geocode each observation to their corresponding DMA.

Table 1.9: Staggered SDiD estimates for all US counties treated with media coverage of East Liverpool event.

	<i>SDiD control treatment specification:</i>	
	All treated DMAs and adjacent counties as controls	All treated DMAs except Youngstown, OH and adjacent counties as controls
Without Covariates	0.571*** (0.180)	0.565*** (0.197)
With Covariates	0.464** (0.202)	0.448** (0.214)
Pre-treatment Mean	3.903	3.892
Observations	13,752	13,656

Note:

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

Treatment variables used here is an indicator for whether a county's DMA has originated media coverage that makes reference to the East Liverpool event. The results for column 1 use a sample for all counties in treated DMAs and their adjacent neighbors, while column 2 excludes counties within the Youngstown OH DMA; both samples drop counties which did not record at least one opioid death per year over the observed period. Covariates employed by the models in row 2 include: Percent of county population hispanic, black, or with a subprime credit score; percent of county laborforce employed in construction; unemployment and poverty rates; annual prescription opioid dispensing rate; arrests for heroin or related drugs per 100 thousand population; and policy indicators for whether the county's state had enacted naloxone access, good samaritan, mandatory PDMP, or opioid prescription restriction laws. Estimates enclosed in brackets are staggered-adoption cluster bootstrap standard errors derived from the method described by Clarke et al. (2023).

level.

1.5.2 Identifying changes to first responder behavior

Understanding how these fentanyl misinformation shocks actually translate to changes in mortality remains an open question. If the mechanism pathway that I have already proposed is valid, than I should be able to examine direct changes to first responder behavior- in particular naloxone administration rates- within the treated regions. Unfortunately, naloxone administration data within Ohio for much of the observed time frame is incomplete. Most notably, naloxone administrations performed by law enforcement are conspicuously absent from the extant data. Additionally, changes to community programs that train and distribute naloxone kits could be affecting their use more than even misinformation.

Though only descriptive, I collect from the Ohio Department of Health data for naloxone

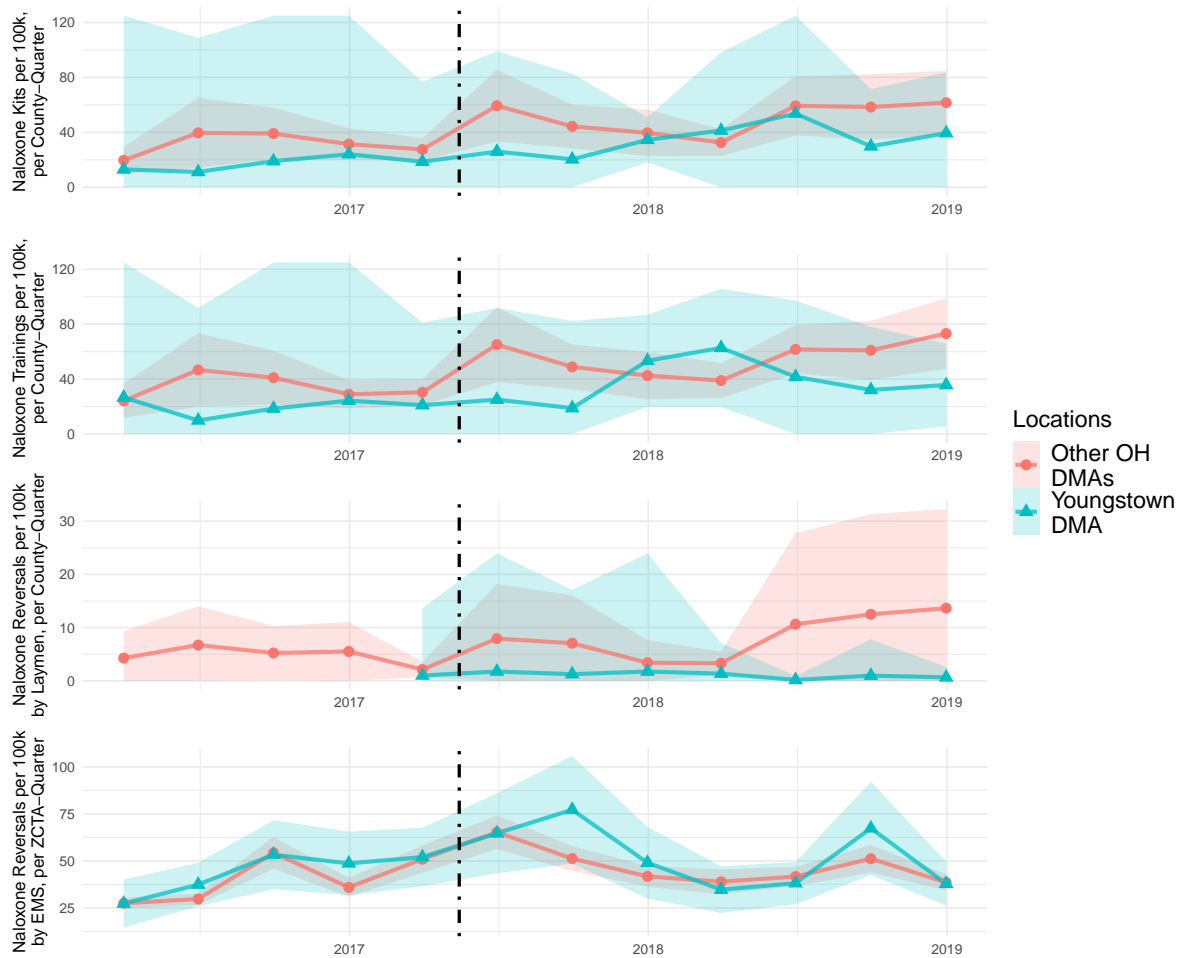
trainings, distributed kits, and self-reported naloxone resuscitations performed by laymen and plot their time series in Figure 1.9. These figures are suggestive only as there is no way to know whether the trained individuals or distributed kits actually remained within the recorded counties. Similarly, resuscitations performed by laymen are almost certainly undercounted. That being said, naloxone distribution and trainings do not appear to significantly differ between the Youngstown OH DMA and the remainder of Ohio counties. Naloxone administrations performed by laymen may be affected by misinformation shocks by staying persistently low relative to the rest of the state, but there is essentially no data for the pre-period (due to data collection not having been initiated until well into 2017 for some counties) with which to make this claim credibly. I additionally collect and plot data for naloxone administrations performed by EMS from the Ohio Department of Public Safety, but similarly find no significant changes between treated and untreated regions. Following observations from Kochersperger (2023) of the relative change in naloxone administration rates in New York state across different first responder types after the East Liverpool media event, it seems likely that if naloxone use is changing substantially among first responders, this is probably most pronounced among unobserved law enforcement officers.

As an alternative approach at concretely describe the underlying behavioral mechanisms that are driving this change in mortality, I examine changes in death locations. If perceptions of the hazards of fentanyl exposure are discouraging the timely administration of aid, then the number of opioid overdose deaths recorded within hospitals would decline in treated relative to untreated areas. I estimate this directly with the following linear probability model (LPM):

$$Y_{c,i,t} = \beta_{11,i}OD_i + \beta_{12,i}Post_t \times OD_i + \beta_{13,i}Post_t \times OD_i \times Treat_c \\ + \delta X_i + \gamma_{c,t} + \varepsilon_{c,i,t}$$

where $Y_{c,i,t}$ is a dichotomous outcome variable representing whether or not individual i , residing in county c that died in month t , has their listed place of death as being in

Figure 1.9: Regional time series of naloxone kit distribution, trainings, and administrations.



((a)) Time series figures for naloxone trainings, distributed kits, self-reported naloxone resuscitations performed by laymen, and naloxone resuscitations performed by EMS per 100k population. Plots are averages across counties/ZCTAs within the Youngstown OH DMA against all other counties in Ohio by county-quarter or ZCTA-quarter. Shaded regions represent the 95% confidence intervals and the dashed vertical line demonstrates the date of the East Liverpool event.

one of five places¹⁸: inpatient hospital setting, outpatient or emergency room hospital setting, dead on arrival to hospital, home or residence, and other. OD_i indicates whether the cause of death is attributed to an opioid overdose, which when interacted with the $Post_t$ and $Treat_c$ dummies imply a triple-differences identification approach. $Post_t$ and $Treat_c$ follow an identical definition to those employed in the staggered-SDiD model based on the Beletsky et al. (2020) data, where the interaction of the two implies a county's DMA has originated media coverage that makes reference to the East Liverpool event. X_i is a vector of individual-level demographic controls, while $\gamma_{c,t}$ are county-month fixed effects. I employ individual death certificate data from the CDC multiple cause of death file and restrict my sample to all deaths attributed to an external injury or poisoning (S00-T88) for the years 2015-2019 within a treated DMA or their adjacent county neighbors. Because opioid-related deaths are relatively uncommon events outside of Appalachia, using alternative external injury deaths as a baseline to compare against would afford a more complete picture first responder practices. As well, this triple-differences approach allows me to observe if there are any structural changes in death locations, beyond just those related to opioid overdoses.

The results of these LPM models are depicted in Table 1.10. To summarize: Within counties that have been exposed to misinformation pertaining to the East Liverpool event, and relative to other causes of death, the likelihood of an opioid overdose death being recorded in an inpatient or outpatient/emergency room setting decreases by approximately 2% and 1.9%, respectively; and the likelihood of those same deaths occurring someplace other than a medical setting or residence increases by approximately 3%. The interpretation of the results within Table 1.10 is that hesitancy in administering aid has yielded fewer attempts at resuscitation, and as such *moved* the location of death from medical to non-medical settings. Considering the urgency of opioid poisoning and the general preventability of death with timely administration of aid, this latter point is troubling.

¹⁸Although I only include five here, the CDC MCODE file includes eight possible values that this location of death variable can take: Hospital, clinic or medical center - inpatient; hospital, clinic or medical center - outpatient or admitted to emergency room; hospital, clinic or medical center - dead on arrival; decedent's home; hospice facility; nursing home/long term care; other; and place unknown. The omitted locations are excluded due to low counts.

Table 1.10: Place-of-death linear probability results.

Death Location:	Inpatient	Outpatient & Emergency Room	DOA	Home	Other
<i>Variables</i>					
Media-Treated	0.0007 (0.0018)	0.0006 (0.0016)	-2.15×10^{-6} (0.0006)	-4.54×10^{-5} (0.0001)	-0.0010 (0.0009)
Is Opioid Overdose	-0.1203*** (0.0018)	-0.0118*** (0.0016)	0.0008 (0.0006)	0.2367*** (0.0023)	-0.0782*** (0.0023)
Media-Treated \times Is Opioid Overdose	-0.0197*** (0.0031)	-0.0186*** (0.0025)	-6.46×10^{-5} (0.0010)	0.0014 (0.0035)	0.0298*** (0.0037)
Pre-treatment Means	0.2609	0.1307	0.0125	0.2720	0.2233
Percent Effect	-0.0755	-0.1423	-0.0051	0.0051	0.1335
<i>Fit statistics</i>					
Observations	809,703	809,703	809,703	809,703	809,703
R ²	0.20116	0.11889	0.11737	0.19103	0.21648
Within R ²	0.01169	0.00050	5.05×10^{-5}	0.03469	0.00363

Note:

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

Clustered (County \times Month) standard-errors in parentheses

Linear probability model estimates for the marginal likelihood of one of five specific death locations being listed on certificate. Treatment defined as the interaction between an indicator for whether an individual death observation's DMA has originated media coverage that makes reference to the East Liverpool event, and an indicator for whether the observed death is attributed to an opioid overdose. Sample is drawn from the complete set of individual deaths recorded in a treated DMA or neighboring county based on Beletsky et al. (2020), which attribute the underlying cause to an external injury or poisoning (corresponding to ICD-10 codes S00-T88). Covariates include reported age, sex, race, ethnicity, highest level of educational attainment, as well as county-month fixed effects.

1.6 Conclusions

In this paper, I have identified the significant impact of fentanyl exposure misinformation on first responder behavior during overdose events and overall opioid-related mortality. In particular, I have examined first responder responses to the widespread dissemination of inaccurate information regarding the supposedly-lethal hazards of passive fentanyl exposure. By analyzing changes in opioid-related mortality following a well-covered episode involving an Ohio police officer in 2017, the study reveals that areas with higher media exposure to this misinformation experience significant increases in opioid overdose deaths. This study underscores the importance of accurate information dissemination and highlights the potentially deadly consequences of misinformation on public health outcomes.

The primary takeaway from my results is that opioid-related mortality appears to be increasing by approximately 2.84 deaths per 100,000 population, per county-quarter; national-scale results place this figure at 0.448 additional deaths per 100,000 population, per county-quarter. For my preferred specification, this jump in mortality represents 56.4% of the average quarterly mortality for the treated counties over the period observed. Back-of-the-envelope calculations identify 199 avoidable overdose deaths, or 72 per year within the Youngstown DMA, according to the SDiD model; and 5,479, or 1,992 per year nationally, according to the preferred staggered-SDiD model. For perspective on these magnitudes in the context of other opioid pandemic policies, Rees et al. (2019) find that the adoption of naloxone access laws by states corresponded to a net decrease of 62-69 opioid-related deaths per year, nationally. Using the Florence et al. (2021) estimate of \$11.548 million in total economic costs per opioid overdose, this would put total costs at \$2.298 and \$48.222 billion within the Youngstown DMA and nationally, respectively.

There are several noteworthy policy implications of my findings. First, it would appear that some corrective effort on the part of criminal justice authorities is needed to combat further dissemination of misinformation. Recent efforts at retraining first responders to correct for fentanyl hazard misperceptions do appear effective (Winograd, Phillips, et al.,

2020; Del Pozo et al., 2021), but there are limitations to the generalizability of these results. An obvious next-step then would be to pursue a randomized control trial experiment to observe the causal influence of first responder retraining on overdose response behavior and mortality.

Enhanced first responder training is an obvious remedy, but there are a number of reasons for policy makers to take pause when considering this particular approach. Namely, narrative correction does not appear to enjoy the same degree of social media play or lurid virality of the initial fentanyl exposure events, so efforts at retraining first responders could be costly if it were required to be conducted at a scale that compensates for this lack of information spillovers; and while there is promise that such retraining can influence first responder beliefs, it is unclear how universally this improved knowledge translates to actual behavior. Analysis of more general overdose education and naloxone distribution training revealed more complicated effects on law enforcement beliefs Winograd, Stringfellow, et al. (2020), with 31% of participants reporting *more* negative attitudes towards overdose victims following training. Similarly, the companion paper to this research (Kochersperger, 2023) observes differing responses to the fentanyl misinformation shock in naloxone administration rates across first responder types, with by far the largest declines in naloxone administration propensities being observed among law enforcement after the East Liverpool media event. This pronounced susceptibility to misinformation points to a broader issue among law enforcement that may be rooted in something more fundamental than a limited knowledge of toxicology. As outlined earlier in the research background, the present fentanyl misinformation panic appears to be just the current iteration of a long-present myth-spinning phenomenon; so even if fentanyl hazards impressions are completely reversed, this may do little to limit future panics.

Another policy consideration is that media coverage of fentanyl hazards need not necessarily promote false information, and can even be a useful means of correcting misperceptions. A cursory survey of recent media coverage of supposed fentanyl exposure incidents does reveal the pronounced use of more skeptical language, and even statements outright dismissing the likelihood of events as described by law enforcement. Still, it is unclear

whether this reactive fact-checking approach is a sufficient means of undoing the damage already wrought by the initial misinformation shocks.

During the early stages of the AIDS epidemic when misbeliefs of the virus' transmission vectors and contagiousness dominated public perceptions, concerted efforts on the part of a select few journalists to correct these narratives were consistently undermined by the broader media environment (Beharrell, 2003). Ultimately, this narrative correction depended on not only the dogged efforts of media insiders, but also celebrity intervention. Noteworthy events such as Princess Diana's visits to the opening of London's Middlesex Hospital AIDS ward in 1987, and Harlem Hospital's AIDS unit in 1989 were seminal in advancing the idea that AIDS victims were not passively contagious after she was photographed shaking hands and hugging them without gloves. Similarly, Earvin "Magic" Johnson's much-publicized HIV-positivity disclosure and subsequent sudden retirement from the NBA has been demonstrated to have reduced stigma surrounding HIV testing, and increased diagnoses among heterosexual men (Cardazzi, Martin, & Rodriguez, 2023). While none of this is to suggest that celebrity endorsements represent a realistic policy response, it does highlight the corrective capacity that media and media consumption can command. Herman et al. (2020) note six months elapsed between the time of the East Liverpool event and the release of the ACMT-AACT joint statement that debunked many of the sensational fentanyl exposure claims; but that over this same time the scientific community was quiet and permitted the unchecked dissemination of misinformation. At a minimum, policymakers should expand on the medical misinformation correcting initiatives pioneered during the COVID-19 pandemic and prioritize swift fact-checking in the future. Effective policy to counter these misinformation narratives must both correct misperceptions among first responders and disincentive the continued dissemination of misinformation by the media.

References

- Abadie, A., Diamond, A., & Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program. *Journal of the American statistical Association*, *105*(490), 493–505.
- Alpert, A., Evans, W. N., Lieber, E. M., & Powell, D. (2022). Origins of the opioid crisis and its enduring impacts. *The Quarterly Journal of Economics*, *137*(2), 1139–1179.
- Alpert, A., Powell, D., & Pacula, R. L. (2018). Supply-side drug policy in the presence of substitutes: Evidence from the introduction of abuse-deterrent opioids. *American Economic Journal: Economic Policy*, *10*(4), 1–35.
- Arias, E., Tejada-Vera, B., Ahmad, F., & Kochanek, K. D. (2022). Vital statistics rapid release: Provisional life expectancy estimates for 2021, report no. 23. *National Center for Health Statistics*.
- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., & Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, *111*(12), 4088–4118.
- Attaway, P. R., Smiley-McDonald, H. M., Davidson, P. J., & Kral, A. H. (2021). Perceived occupational risk of fentanyl exposure among law enforcement. *International Journal of Drug Policy*, *95*, 103303.
- Beharrell, P. (2003). Aids and the british press. In *Getting the message* (pp. 202–238). Routledge.
- Beletsky, L., Seymour, S., Kang, S., Siegel, Z., Sinha, M. S., Marino, R., . . . Freifeld, C. (2020). Fentanyl panic goes viral: The spread of misinformation about overdose risk from casual contact with fentanyl in mainstream and social media. *International Journal of Drug Policy*, *86*, 102951.
- Berardi, L., Bucerius, S., Haggerty, K. D., & Krahn, H. (2021). Narcan and narcan't: Implementation factors influencing police officer use of narcan. *Social science & medicine*, *270*, 113669.
- Bonisteel, S. (2015, May). *Utah foots the bill for ailing cops' controversial*

- scientology-based detox treatment*. FOX News Network. Retrieved from <https://www.foxnews.com/story/utah-foots-the-bill-for-ailing-cops-controversial-scientology-based-detox-treatment>
- Bucerius, S., Berardi, L., Haggerty, K. D., & Krahn, H. (2022). New drugs, new fears: synthetic opioids and adaptations to police practice. *Policing and Society*, *32*(7), 832–845.
- Buchanich, J. M., Balmert, L. C., Williams, K. E., & Burke, D. S. (2018). The effect of incomplete death certificates on estimates of unintentional opioid-related overdose deaths in the united states, 1999-2015. *Public Health Reports*, *133*(4), 423–431.
- Bump, B. (2018). *How dangerous is fentanyl to first responders? Confusion abounds as use soars* — *timesunion.com*. <https://www.timesunion.com/news/article/How-dangerous-is-fentanyl-Confusion-abounds-as-13224686.php>. ([Accessed August, 2023])
- Cardazzi, A., Martin, J. C., & Rodriguez, Z. (2023). Information shocks and celebrity exposure: The effect of “magic” johnson on aids diagnoses and mortality in the us. *Health Economics*.
- Carrieri, V., Madio, L., & Principe, F. (2019). Vaccine hesitancy and (fake) news: Quasi-experimental evidence from italy. *Health economics*, *28*(11), 1377–1382.
- Chou, W.-Y. S., Oh, A., & Klein, W. M. (2018). Addressing health-related misinformation on social media. *Jama*, *320*(23), 2417–2418.
- Ciccarone MD, D. (2020). No moral panic: Public health responses to illicit fentanyls. *DePaul Law Review*, *70*(2), 229.
- Clarke, D., PailaÑir, D., Athey, S., & Imbens, G. (2023). Synthetic difference in differences estimation. *arXiv preprint arXiv:2301.11859*.
- Conley, T. G., & Taber, C. R. (2011). Inference with “difference in differences” with a small number of policy changes. *The Review of Economics and Statistics*, *93*(1), 113–125.
- DEA. (2021). National forensic laboratory information system: Nflisdrug 2020 annual report.

- Del Pozo, B., Sights, E., Kang, S., Goulka, J., Ray, B., & Beletsky, L. A. (2021). Can touch this: training to correct police officer beliefs about overdose from incidental contact with fentanyl. *Health & Justice, 9*(1), 1–6.
- Devitt, M. (2018). Cdc data show us life expectancy continues to decline. *Am Fam Physician.*
- Doleac, J. L., & Mukherjee, A. (2022). The effects of naloxone access laws on opioid abuse, mortality, and crime. *The Journal of Law and Economics, 65*(2), 211–238.
- Erfanian, E., Grossman, D., & Collins, A. R. (2019). The impact of naloxone access laws on opioid overdose deaths in the us. *Review of Regional Studies, 49*(1), 45–72.
- Feldman, R., & Weston, B. W. (2022). Accidental occupational exposure to a large volume of liquid fentanyl on a compromised skin barrier with no resultant effect. *Prehospital and disaster medicine, 37*(4), 550–552.
- Florence, C., Luo, F., & Rice, K. (2021). The economic burden of opioid use disorder and fatal opioid overdose in the united states, 2017. *Drug and alcohol dependence, 218*, 108350.
- Fowler, C. S., & Jensen, L. (2020). Bridging the gap between geographic concept and the data we have: The case of labor markets in the usa. *Environment and Planning A: Economy and Space, 52*(7), 1395–1414.
- Greenwood, L., & Fashola, K. (2021). *Illicit fentanyl from china: An evolving global operation*. US-China Economic and Security Review Commission.
- Hedegaard, H., Miniño, A. M., Spencer, M. R., & Warner, M. (2021). Drug overdose deaths in the united states, 1999–2020.
- Herman, P. A., Brenner, D. S., Dandorf, S., Kemp, S., Kroll, B., Trebach, J., ... Stolbach, A. I. (2020). Media reports of unintentional opioid exposure of public safety first responders in north america. *Journal of Medical Toxicology, 16*(2), 112–115.
- Jones, A. E. (2023, Apr). *Claim that fentanyl is leading cause of death for u.s. adults 18 to 45 needs context*. Retrieved from <https://www.verifythis.com/article/news/verify/national-verify/fentanyl-leading-cause-of-death-for-american-adults-age-18-45-claim-needs-context/536-555c1f22-b1bf>

- 40cc-b9bc-6ee9ce822b71
- Kaul, A., Klößner, S., Pfeifer, G., & Schieler, M. (2022). Standard synthetic control methods: The case of using all preintervention outcomes together with covariates. *Journal of Business & Economic Statistics*, *40*(3), 1362–1376.
- Kochersperger, E. (2023). Fentanyl on My Mind: Perceived Opioid Exposure Risk and its Influence on Naloxone Administration Rates. *Working Paper*.
- Kranz, S. (2022). *Synthetic difference-in-differences with time-varying covariates* (Tech. Rep.). Technical report. Available online at: <https://github.com/skranz/xsynthdid>
- Lynch, M. J., Suyama, J., & Guyette, F. X. (2018). Scene safety and force protection in the era of ultra-potent opioids. *Prehospital emergency care*, *22*(2), 157–162.
- Macmadu, A., Yolken, A., Frueh, L., Toussaint, J. R., Newman, R., Jacka, B. P., . . . Marshall, B. D. (2022). Characteristics of events in which police responded to overdoses: an examination of incident reports in rhode island. *Harm Reduction Journal*, *19*(1), 116.
- Moss, M. J., Warrick, B. J., Nelson, L. S., McKay, C. A., Dubé, P.-A., Gosselin, S., . . . Stolbach, A. I. (2018). Acmt and aact position statement: preventing occupational fentanyl and fentanyl analog exposure to emergency responders. *Clinical Toxicology*, *56*(4), 297–300.
- Packham, A. (2019). *Are syringe exchange programs helpful or harmful? new evidence in the wake of the opioid epidemic* (Tech. Rep.). National Bureau of Economic Research.
- Persaud, E., & Jennings, C. R. (2020). Pilot study on risk perceptions and knowledge of fentanyl exposure among new york state first responders. *Disaster medicine and public health preparedness*, *14*(4), 437–441.
- Peterson, A. B., Gladden, R. M., Delcher, C., Spies, E., Garcia-Williams, A., Wang, Y., . . . others (2016). Increases in fentanyl-related overdose deaths—florida and ohio, 2013–2015. *Morbidity and Mortality Weekly Report*, *65*(33), 844–849.
- Pickett, R. E., Hill, J., & Cowan, S. K. (2022). The myths of synthetic control:

- Recommendations for practice.
- Ramos-Matos, C. F., Bistas, K. G., & Lopez-Ojeda, W. (2022). Fentanyl. In *Statpearls [internet]*. StatPearls Publishing.
- Rees, D. I., Sabia, J. J., Argys, L. M., Dave, D., & Latshaw, J. (2019). With a little help from my friends: the effects of good samaritan and naloxone access laws on opioid-related deaths. *The Journal of Law and Economics*, *62*(1), 1–27.
- Riches, J. R., Read, R. W., Black, R. M., Cooper, N. J., & Timperley, C. M. (2012). Analysis of clothing and urine from moscow theatre siege casualties reveals carfentanil and remifentanil use. *Journal of analytical toxicology*, *36*(9), 647–656.
- Scientology Critical Information Directory. (2009). *Bio-cleansing centers of america, llc*. Retrieved from <http://www.xenu-directory.net/documents/corporate/entity.php?ntt=292>
- Smart, R., Pardo, B., & Davis, C. S. (2021). Systematic review of the emerging literature on the effectiveness of naloxone access laws in the united states. *Addiction*, *116*(1), 6–17.
- Smiley-McDonald, H. M., Attaway, P. R., Richardson, N. J., Davidson, P. J., & Kral, A. H. (2022). Perspectives from law enforcement officers who respond to overdose calls for service and administer naloxone. *Health & Justice*, *10*(1), 9.
- Stanley, T. H. (1992). The history and development of the fentanyl series. *Journal of pain and symptom management*, *7*(3), S3–S7.
- Stoecker, W. V., Madsen, D. E., Cole, J. G., & Woolsey, Z. (2016). Boys at risk: Fatal accidental fentanyl ingestions in children: Analysis of cases reported to the fda 2004–2013. *Missouri medicine*, *113*(6), 476.
- Wang, Y., McKee, M., Torbica, A., & Stuckler, D. (2019). Systematic literature review on the spread of health-related misinformation on social media. *Social science & medicine*, *240*, 112552.
- White, M. D., Watts, S., Orosco, C., Perrone, D., & Malm, A. (2022). Leveraging body-worn camera footage to better understand opioid overdoses and the impact of police-administered naloxone. *American journal of public health*, *112*(9), 1326–1332.

- Winograd, R. P., Phillips, S., Wood, C. A., Green, L., Costerison, B., Goulka, J., & Beletsky, L. (2020). Training to reduce emergency responders' perceived overdose risk from contact with fentanyl: early evidence of success. *Harm Reduction Journal*, *17*(1), 1–5.
- Winograd, R. P., Stringfellow, E. J., Phillips, S. K., & Wood, C. A. (2020). Some law enforcement officers' negative attitudes toward overdose victims are exacerbated following overdose education training. *The American journal of drug and alcohol abuse*, *46*(5), 577–588.
- Zibbell, J. E., Clarke, S. D., Kral, A. H., Richardson, N. J., Cauchon, D., & Aldridge, A. (2022). Association between law enforcement seizures of illicit drugs and drug overdose deaths involving cocaine and methamphetamine, ohio, 2014–2019. *Drug and alcohol dependence*, *232*, 109341.

2 Fentanyl on My Mind: Perceived Opioid Exposure Risk and its Influence on Naloxone Administration Rates

2.1 Introduction

The opioid epidemic is one of the most pressing public health crisis of our era, and has seen only a marked exacerbation since the onset of the COVID-19 pandemic, further complicating an already dire situation. Conventional approaches to combating this epidemic, including interdiction efforts aimed at curtailing the flow of illicit drugs and educational campaigns directed at vulnerable populations, have largely proven ineffectual in stemming the tide of opioid-related mortality. In an attempt to mitigate this crisis, policymakers have turned towards innovative harm reduction strategies designed to diminish the risks faced by opioid users. Among these, Good Samaritan laws, syringe service programs, and fentanyl test strip access initiatives stand out for their dual objectives: Reducing the legal repercussions for individuals assisting overdose victims and decreasing the health risks associated with opioid usage. Proponents of these policies argue that they do not promote drug usage but rather provide individuals suffering from substance use disorders with the resources necessary to make informed decisions and embark on a path towards recovery.

Naloxone access laws (NALs) represent a pivotal aspect of these harm reduction efforts, aiming to broaden the availability of the opioid antagonist—commercially known as Narcan—and lessen the legal barriers to its use. Administered in a timely manner, naloxone has the capacity to revive individuals from overdose-induced respiratory failure, thereby presenting a critical lifeline. The objective behind NALs is to curtail opioid-related fatalities by enhancing the accessibility of this vital medication. However, empirical evaluations of these laws paint a complex picture. While NALs have been generally effective in reducing opioid mortality, they have also been critiqued for potentially fostering *ex ante* moral hazard, whereby the diminished risk of overdose death might inadvertently encourage

continued opioid consumption. To date, research has predominantly concentrated on quantifying the magnitude of these moral hazard effects (Doleac & Mukherjee, 2018), seeking to ascertain whether they merely attenuate or entirely negate the benefits conferred by naloxone.

This study delves into a nuanced aspect of the opioid crisis by examining how misinformation concerning the risks associated with administering aid influences individuals' willingness to assist in overdose situations. The act of delivering naloxone necessitates close physical proximity between the rescuer and the overdose victim, raising concerns about potential harm to the rescuer. This apprehension is exacerbated by widespread, albeit unfounded beliefs among first responders that even momentary contact with fentanyl—a potent opioid—could prove fatal. Such misconceptions, largely propagated by a few high-profile media incidents coinciding with the widespread adoption of NALs, suggest an alternative behavioral dynamic that may explain the mixed efficacy of these laws. The increased availability of naloxone, without a corresponding decrease in mortality rates, hints at a reluctance among potential rescuers to utilize this critical intervention tool due to fear of personal harm, thus undermining the effectiveness of NALs in combating the opioid epidemic.

By exploiting a unique dataset on naloxone administrations that distinguishes between first responder types from the New York State Department of Health, this analysis examines the influence of media-fueled misinformation shocks on first responder behavior during opioid overdose episodes. I concentrate on a 2017 incident involving the hospitalization of a law enforcement officer after he was exposed to a substance believed to be fentanyl. This event, as highlighted in media analyses like Beletsky et al. (2020) study, was a key factor in entrenching the fallacious idea that fentanyl is so uniquely dangerous it poses lethal acute health risks through momentary contact in the collective public mindset. Leveraging the unanticipated timing of this event, I seek to explore how misinformation dissemination influences opioid-related mortality rates. Through a differences-in-differences (DiD) analysis that takes advantage of the distinct levels of awareness regarding the hazards of opioid exposure (essentially, how susceptible they are to this misinformation) between

harm reduction professionals and other emergency responders, my findings point to a marked decrease in the number of naloxone administrations by police officers following the dissemination of misinformation. This decrease is interpreted as a sign of increased caution among officers, likely due to concerns over the amplified risks fentanyl is perceived to present.

This study can be viewed as a companion piece to Kochersperger (2023), in that these papers are providing the first credibly-causal estimates for the economic consequences of the fentanyl hazards myth. Kochersperger (2023) examines changes in county-level opioid-related mortality rates, leveraging variation in coverage of the same 2017 misinformation employed here across media markets. While Kochersperger (2023) examines down-stream opioid mortality responses to media coverage surrounding the fentanyl misinformation panic, this paper focuses on the underlying mechanism of first responder and naloxone administration. Despite a wealth of qualitative analysis into how misinformation might shape first responder behaviors (Attaway, Smiley-McDonald, Davidson, & Kral, 2021; Beletsky et al., 2020; Del Pozo et al., 2021; Herman et al., 2020), it is only these two studies that have looked into the tangible effects these attitudes have on public health indicators. This study contributes a new dimension to understanding the complexity behind the effectiveness of NALs by highlighting the critical influence of misinformation on public health initiatives. It moves beyond the conventional moral hazard debate, proposing that the reluctance of law enforcement to administer naloxone, even when it is readily accessible, might explain the inconsistent empirical results. This insight marks a unique addition to the economic discourse surrounding NALs and the broader framework of harm reduction strategies.

The remainder of the paper is organized as follows. In the next section I offer some context on naloxone access laws and the specific initiatives employed in New York State to improve naloxone availability, as well as a brief summary of the fentanyl hazards myth and relationship to law enforcement. In Section 3 I describe my data and empirical strategy, in Section 4 I report my results, and in Section 5 I offer a discussion of the policy implications. I conclude in Section 6.

2.2 Background

2.2.1 Naloxone and New York state

Harm reduction policies have taken one of two common forms: Efforts aimed at encouraging overdose victims or their peers to more readily seek medical assistance when in need, and efforts that lower the direct health risks of drug use. The former of these policies are perhaps best typified by Good Samaritan laws, which reduce or eliminate criminal penalties for drug use and possession by parties that seek emergency medical aid for overdosing peers. The latter has been historically associated with syringe service programs (or needle exchanges) that furnish intravenous drug users (IDU) clean syringes to limit the spread of blood-borne disease. Lately however, there have been increased efforts by policymakers to improve access to the important opioid antagonist, naloxone. While naloxone first received approval for treating opioid use disorder in 1971, its use within emergency settings to promptly reverse respiratory failure from acute opioid toxicity has expanded considerably with the development of nasal aspirator and autoinjector delivery methods. Related harm reduction initiatives, naloxone access laws (NALs), have involved a combination of loosening prescribing requirements and limiting liability that bystanders face for administering naloxone, as well as directly improving the availability of the lifesaving drug through kit distribution campaigns.

Despite its well demonstrated life-saving capabilities, a common moral hazard concern with expanding access to naloxone is that it might induce more reckless substance use behavior by reducing the perceived risks of overdose. The concern that such policies might ‘enable’ additional drug usage, echoing the objections to clean needle exchanges dating as far back as the 1980s, has been a common refrain among policymakers invoking this moral hazard argument. The recent endeavor of researchers has been to quantify the degree to which the life-saving advantages of naloxone are undermined by moral hazard concerns. These empirical efforts have mostly focused on examining changes to opioid-related mortality following adoption of NALs. A systematic review of this literature

found that NALs demonstrate slight declines to opioid-related mortality (Smart, Pardo, & Davis, 2021), but that these results are hardly consistent. For instance, McClellan et al. (2018); Rees et al. (2019) look at NAL adoption while also accounting for other simultaneous harm reduction policies and find that improved naloxone access yields significant declines to opioid mortality. Conversely, Doleac and Mukherjee (2018) find evidence of substantial moral hazard effects and large net *increases* to mortality with the adoption of NALs. One concern with this line of research is that NAL adoption may not be sufficiently random to permit a causal interpretation of these results. Other research efforts have explicitly attempted to model for states' selection into treatment, and when doing so find no significant change to opioid mortality (Erfanian, Grossman, & Collins, 2019). Nonetheless, these remain a popular if ambiguously-effective policy within the harm reduction toolkit, with all 50 states and the District of Columbia having adopted some form of NAL by the end of 2017¹⁹. One motivation for performing this research here is to examine whether a contemporaneous shock to first responder behavior may be discouraging naloxone's use, even as its availability has improved.

New York State has been at the forefront of integrating harm reduction policies to combat the opioid epidemic, with legislative measures dating back to 2006 aimed at empowering potential witnesses to opioid overdoses through naloxone training programs registered with the New York State Department of Health (NYSDOH). The introduction of the 911 Good Samaritan Law in September 2011, aimed at alleviating fears of legal repercussions when reporting overdoses, along with the approval for the use of mucosal atomizer naloxone devices by Basic Life EMS agencies in 2013, significantly broadened the scope of naloxone's accessibility. Furthermore, the establishment of Opioid Overdose Prevention Programs (OOPP) in 2014 marked a critical step in training non-medical persons for opioid overdose management and enhancing public access to naloxone. These initiatives underscore New York's commitment to reducing opioid mortality through legislative and regulatory avenues, and key to the research design employed here, all policies were implemented prior to the period analyzed in this study.

¹⁹See Prescription Drug Abuse Policy System (PDAPS), <https://pdaps.org/datasets/laws-regulating-administration-of-naloxone-1501695139>.

In the realm of law enforcement, the 2014 regulations further permitted “shared access” to naloxone, which allows officers to access naloxone within first aid kits under a non-patient-specific prescription. This change, coupled with the commencement of the statewide law enforcement naloxone initiative in the same year, expanded naloxone availability and overdose scenario training among law enforcement officers, recognizing them as often the first responders to opioid overdoses. The motivation for this is clear: Law enforcement are often the first line of defense when responding to opioid overdoses. Empirical evidence, including officer surveys and analysis of bodycam footage, underscores the pivotal role of law enforcement in immediate overdose response, often beating EMS to the scene (Smiley-McDonald, Attaway, Richardson, Davidson, & Kral, 2022; White, Watts, Orosco, Perrone, & Malm, 2022). Within New York state between 2015 and 2020, law enforcement personnel were the first on scene in 85% of cases in which they administered naloxone (Pourtaher et al., 2022); though, in 37% of these cases at least one other dose is administered by some other party (either a layperson or EMS). This is in line with broader naloxone protocols: LEOs often immediately administer naloxone then attempt CPR, so as to buy time before EMS arrives, who ultimately may administer additional doses as needed (Smiley-McDonald et al., 2022). Indeed, state law dictates that all naloxone-trained parties (both LEO and laypersons) notify EMS immediately when responding to a scene, regardless of intervention outcomes prior to EMS arrival. This can be readily observed by the fact that in 91.2% of instances which LEOs responded to, care is transferred to EMS and victims are transported to a hospital (Pourtaher et al., 2022). If not escorted to a hospital, surviving victims are either arrested or released at the scene. Regardless of these post-intervention outcomes however, State law requires all LEOs to submit a timely incident report that outlines their naloxone use.

Community Opioid Overdose Prevention programs (COOP) facilitate both the training of laypersons in proper naloxone use, and the distribution of naloxone kits to eligible parties. These programs have existed in some form or another in the state since 2006, but the 2014 law updated naloxone distribution, and intervention reporting and training procedures. The broader program traces its roots to the State’s network of 25 syringe exchange programs

and adopts an explicitly harm reduction-oriented approach to combating the opioid epidemic. Registries from 2019²⁰ show that the majority of these non-law enforcement opioid prevention programs are composed of local health departments, drug treatment or support, and outpatient medical facilities. This aspect is important, as it highlights the typical familiarity of the COOP-affiliated individuals who are performing these overdose reversals with narcotics exposure hazards. Despite the structured framework for naloxone training and distribution, there exists a notable gap in data regarding the frequency with which laypersons, as opposed to EMS or LEOs, are first responders to overdose scenes. The geographical and operational positioning of COOPs, often in proximity to areas frequented by opioid users, suggests a theoretical advantage in facilitating quicker overdose interventions by trained laypersons. This hypothesis is somewhat supported by statistics from syringe access programs within New York state, which reported 68% of naloxone administrations were carried out by laypersons (Pourtaher et al., 2022). While state law mandates that all trained individuals submit naloxone administration reports through their affiliated COOP, these procedures are not as rigorously adhered to as with LEOs.

2.2.2 The fentanyl misinformation panic

A critical aspect of the opioid epidemic of late is the unprecedented influence of the opioid fentanyl. From 2013 to 2020, deaths linked to synthetic opioids (which are predominantly composed of fentanyl and its close analogs) surged by 18 times, making up 82% of all opioid-related fatalities in 2020 (Hedegaard, Miniño, Spencer, & Warner, 2021). This dramatic increase in fentanyl's market presence has spurred considerable interest and discussion on the hazards associated with incidental exposure to the substance. Analysis performed by the American College of Medical Toxicology and American Academy of Clinical Toxicology has effectively dismissed concerns that short-term contact with fentanyl is dangerously toxic (Moss et al., 2018), revealing that only exposure lasting several hours could lead to significant absorption either through inhalation or skin contact. Contrary to these findings however, media reports often sensationalize encounters, especially with

²⁰See archived regional registries here:
https://www.health.ny.gov/diseases/aids/general/resources/oop_directory/index.htm

law enforcement, as having immediate, severe health consequences from minor exposures, implying fentanyl's involvement. Additional analyses have not substantiated any cases of poisoning from these reported exposures (Lynch, Suyama, & Guyette, 2018; Herman et al., 2020; White et al., 2022), suggesting that the symptoms are more likely attributable to stress-induced psychosomatic responses rather than direct chemical harm.

Despite doubts of their accuracy, reports of first responders' experiences with fentanyl feature prominently in the media. Influenced by this media coverage, a significant part of the first responder community, including law enforcement and EMS workers, appears to believe in the severe danger posed by brief contact with fentanyl, as evidenced by high agreement rates in surveys (Persaud & Jennings, 2020; Del Pozo et al., 2021; Attaway et al., 2021; Berardi, Bucerus, Haggerty, & Krahn, 2021; Bucerus, Berardi, Haggerty, & Krahn, 2022). These misconceptions are consequential, influencing some first responders to avoid providing aid for fear of fentanyl exposure (Berardi et al., 2021; Bucerus et al., 2022).

Beletsky et al. (2020) highlighted a key incident in East Liverpool, OH, in May 2017, which had a profound impact on the collective anxiety surrounding fentanyl exposure. During this incident, a police officer reportedly came into contact with fentanyl while making an arrest, leading to severe symptoms that necessitated hospitalization after multiple administrations of naloxone. The widespread media attention and social media dissemination, especially on Facebook, were instrumental in intensifying public concern over fentanyl. Analysis of contemporaneous Google search trends highlights a surge in fentanyl exposure interest following the 2017 event (Kochersperger, 2023). Kochersperger (2023) further explores the effect of the media's portrayal of this incident on opioid-related deaths, utilizing the uneven coverage across different regions to assess its impact. Their findings suggest a significant correlation between the media focus on this event and an increase in opioid-related fatalities. They propose that this surge in deaths may be linked to first responders' reluctance to engage in overdose interventions, fearing fentanyl exposure. However, while their research effectively connects media coverage with rising mortality rates, it stops short of pinpointing the exact mechanism. In this analysis, I aim

to build upon the East Liverpool event identification approach, enhancing it with data that more directly accounts for first responder Naloxone use when responding to opioid overdoses.

A pivotal element in dissecting the public's fear regarding fentanyl relates to the varied perceptions among first responders. Bucerius et al. (2022) and Kochersperger (2023) describe LEOs' historical vulnerability to misinformation, particularly during health crises amplified by media narratives. This vulnerability is accentuated in the fentanyl discourse, where LEOs' exposure to misinformation is compounded by the nature of media-reported exposure incidents, predominantly featuring law enforcement personnel. Furthermore, certain training protocols for LEOs have officially propagated the notion of fentanyl exposure risks (Kucher & Figueroa, 2021). In contrast, harm reduction advocates and medical professionals have consistently endeavored to rectify this misleading narrative. Despite the absence of comprehensive surveys to measure the prevalence of the fentanyl hazard myth within medical and harm reduction communities, anecdotal accounts reveal a widespread skepticism towards the accuracy of media portrayals ²¹.

Considering that many of the laymen being trained by these experts (and who consequently constitute a portion of the observed COOP resuscitations within the data) are themselves either current or past opioid users (NYSDOH, 2023), intimate personal knowledge and experience in handling illicit opioids would preclude a susceptibility to the exposure hazards misinformation²². While there is evidence indicating that users do change their consumption practices in the presence of fentanyl (Rouhani, Park, Morales, Green, &

²¹Representatives of the National Harm Reduction Coalition aired their frustration in 2018 with having the combat the fentanyl "bioterrorism threat" narrative being promoted by media: <https://harmreduction.org/blog/fentanyl-exposure/>. Also in 2018, other harm reduction advocates employed direct methods to combat this misinformation by demonstrating the benign nature of powdered fentanyl: <https://twitter.com/chadsabora/status/1024440190889287680>. The AMA adopted policy in 2019 which formalized their position of combatting erroneous beliefs surrounding the fentanyl hazards myth, citing a particular concern that first responders may be more reluctant to utilize naloxone in life threatening situations: <https://www.ama-assn.org/system/files/2019-05/a19-yps-resolution-02.pdf>

²²For instance, self-reported aversion to, and even phobia of needles is relatively common among IDUs (McBride, Pates, Arnold, & Ball, 2001; Tompkins, Ghoneim, Wright, Sheard, & Jones, 2007). Because needle-phobic IDUs willingly continue to use fentanyl through injection however (and user surveys affirm that injection is by far the most common route employed (Buresh, Genberg, Astemborski, Kirk, & Mehta, 2019)), this would suggest an awareness of the implausibility of receiving a high through presumably preferred, passive routes.

Sherman, 2019), there is nothing to suggest that they outright avoid the opioid out of some concern for exposure hazards. The vulnerability of EMS to misinformation is more opaque. While there are supposed exposure incidents involving EMS first responders, these are considerably less common as those involving LEOs (Adams, Maloy, & Warrick, 2023). Still, surveys of New York state EMS do demonstrate that the majority hold many of the fictitious beliefs associated with the fentanyl hazards myth, though with slightly lower prevalence compared to LEOs (Adams et al., 2023).

2.3 Empirical approach

2.3.1 Data

Since the principal outcome variable of interest in this analysis is overdose resuscitation, an ideal set of data would include individual event-level data for every instance that naloxone was or could have feasibly been administered, including information on the time and precise location of the event, as well as details on the first responder; for the treatment variable, one would have information on the personal beliefs of each first responder involved in every overdose episode. While the latter of these two is almost certainly nonexistent, the former does have several close approximations. Numerous states report monthly naloxone administration data to the county level, but a cursory investigation reveals that much of this data is limited in availability to only events occurring between 2018-present.

For the primary component of my analysis here I collect data from the New York State Department of Health (NYSDOH) on the total number of unique Naloxone administration events by first responder type and by county-quarter over the period Q1 2015-Q4 2019. These data correspond to the state's 57 "upstate" and Long Island counties and exclude New York City. EMS data is derived from the NEMESIS system, to which the vast majority of EMS agencies report; state law requires all Naloxone resuscitation attempts performed by either law enforcement officers (LEO) and community opioid overdose prevention (COOP) programs to be reported to the NYSDOH. These data also include the non-suppressed county-quarter overdose death counts. Summary statistics for this data can be

Table 2.1: Summary Statistics of New York Data by County-Quarter

Statistic	N	Mean	Median	St. Dev.
LEO Naloxone Administrations	1,140	5.739	2	11.168
EMS Naloxone Administrations	1,140	29.675	12.5	42.333
COOP Naloxone Administrations	1,140	5.697	1	13.829
Deaths (All Opioids)	1,140	7.582	2	14.233
Deaths (Heroin)	1,140	2.837	1	5.287
Deaths (Other Opioids)	1,140	6.386	2	12.715
ED Visits (All Opioids)	798	36.058	16	57.479
ED Visits (Heroin)	706	28.076	13	43.618
ED Visits (Other Opioids)	461	15.217	8	21.326
Hospitalizations (All Opioids)	533	13.741	8	19.063
Hospitalizations (Heroin)	607	3.857	0	8.183
Hospitalizations (Other Opioids)	498	8.004	0	12.064
Rehab Admissions (Heroin)	935	197.043	82	287.369
Rehab Admissions (Other Opioids)	935	63.618	38	83.722
% Pop. Sub-prime Credit	1,140	23.537	23.506	3.432
% Pop. Medicaid	1,140	8.538	7.611	3.549
Unemployment Rate	1,140	4.915	4.700	1.186
% Labor Force in Construction	1,140	0.042	0.040	0.015
Average HH Earnings	1,140	3,899.339	3,754	666.422

All data described above represents average observed values by county-quarter for the years 2015-2019. Data for the total number of unique Naloxone administration events by first responder type and non-suppressed county-quarter overdose death counts by opioid type are from the New York State Department of Health (NYSDOH). Counts for the number of opioid overdose-related emergency department (ED) visits and inpatient hospitalizations come from the New York Statewide Planning and Research Cooperative (SPARCS) and unique drug treatment clinic admissions from the Office of Addiction Services and Supports (OASAS) systems. Other data for the percentage of the population with a credit score below 660 is collected from Equifax, percent of population eligible for Medicaid from the NYSDOH, unemployment rate, percent of laborforce employed in construction, and average household earning are from the BLS quarterly census of employment and wages.

found in Table 3.1. In Figure 2.1 I plot average county-quarter opioid mortality figures over time. Without any additional context, one can already see that the timing of the East Liverpool event (represented by the dashed vertical line) appears to correspond to a sizeable change in mortality patterns.

In addition to these mortality and naloxone administration data, I also collect opioid-related hospitalization and drug rehabilitation program admissions data from the same public-use NYSDOH source. These data originate from the Statewide Planning and Research Cooperative (SPARCS) and Office of Addiction Services and Supports (OASAS) systems, respectively, but unlike the mortality data are suppressed over certain county-quarters²³. These data include county-quarter counts for the number of opioid overdose-related emergency department (ED) visits, inpatient hospitalizations, and unique drug treatment clinic admissions²⁴. To account for any confounding effects, I also collect county-level quarterly estimates for the percent of population with a sub-prime credit score from Equifax, percent of population eligible for Medicaid from the NYSDOH, unemployment rate, percent of laborforce employed in construction, and average household earning from the BLS quarterly census of employment and wages.

2.3.2 Identification strategy


Media analyses (Beletsky et al., 2020) have identified a particularly noteworthy 2017 event in East Liverpool, Ohio²⁵ as being an early major influence in the proliferation of fentanyl hazards misinformation. For identification purposes, I utilize the random timing of this event in a differences-in-differences setting to identify the causal treatment of misinformation on first responder behavior. An effective approach to achieve this is

²³To avoid patient identification, SPARCS hospitalization data is suppressed when a county-quarter records 1-5 cases, and OASAS data is suppressed when a county-quarter records 1-10 cases. In practice, this translates to a slight over representation of larger, more urban counties. As well, since opioid use has increased significantly over the timespan examined, these hospitalization and rehabilitation admissions data exhibit greater coverage in the later time periods.

²⁴Starting January 2020, NYSDOH redefined the reported treatment program figures as based on the number of admissions during the quarter or year, and not on the number of individuals admitted or individuals treated. Due to these changes in reporting practices, comparable OASAS data is only available for as recent as Q1 2019 since reporting of this data is delayed by several months.

²⁵see: <https://www.wcpo.com/news/state/state-ohio/police-east-liverpool-officer-accidentally-overdoses-on-fentanyl-after-making-drug-arrest>

Figure 2.1: Average opioid-related death rates and counts timeseries county-quarter and opioid type.



Average opioid-related deaths per 100 thousand population, and opioid-related death counts by opioid type per county-quarter. Shaded regions represent the 95% confidence intervals and the dashed vertical line demonstrates the date of the East Liverpool event. The term "pain relievers" encompasses fatalities from overdoses related to both pharmaceutical and illicitly manufactured opioids, including substances like fentanyl.

to delve deeper into the qualitative distinctions in how misinformation is spread and perceived among different types of first responders. In particular, I examine the relative change in naloxone administration between LEOs, EMS, and COOP programs following the East Liverpool event. It is worth noting that the vast majority of the exposure events reported on within the media have been pertaining to events involving LEOs, while nary few have involved EMS and seemingly none COOPs. Considering that most of the research on this phenomenon and de-bunking of exaggerated exposure claims has been carried out by harm reduction researchers and advocates, this is unsurprising. It is plausible to infer that first responders associated with a COOP are likely less swayed by media depictions of fentanyl exposure risks due to their more extensive experience with drug overdoses and their existing understanding of the minimal tactile dangers associated with fentanyl.

A conventional differences-in-differences design should be able to identify misinformation effects on first responders by separating LEOs and/or EMS as the treated and COOPs

as the control and observing changes in naloxone administration before and after the liverpool event. To do so, I estimate the following regression:

$$Y_{c,i,t} = \beta_1 \text{Post}_t + \beta_{2,i} \text{Treat}_i + \beta_{12} \text{Post}_t \times \text{Treat}_i + \delta X_{c,t} + \gamma_{i,t} + \mu_c + \eta_i + \lambda_t + \varepsilon_{c,i,t}, \quad (2.1)$$

where $Y_{c,i,t}$ is the total number of episodes where naloxone is administered in county c , by first responder type i , in time t ; Treat_i is an indicator for whether the type of first responder is considered to be treated (I estimate two different specifications where the treated are considered to be only LEOs, and another where both LEOs and EMS are treated); and $X_{c,t}$ is a vector of local socio-economic controls.

The validity of this approach is contingent on the parallel trends assumption. To demonstrate the existence of parallel pre-treatment trends, in Figure 2.2 I plot average county-quarter naloxone administrations by first responder type. While EMS naloxone administrations appear to be exhibiting some unusual pre-trends, LEO and COOP behavior follow similar patterns until the East Liverpool event— suggesting that DiD is an appropriate design here. I explicitly assume the incident in East Liverpool did not cause a differential change in the volume of calls to various first responders or healthcare providers. This assumption holds weight for several reasons: East Liverpool’s geographical location in another state makes local spillover effects improbable; there were no simultaneous changes to the COOP program’s structure or overall harm reduction policies in New York state; and individuals in need of emergency care do not have the agency to select which first responder attends to their emergency, indicating that any observed changes in administration rates are more indicative of responder behavior than that of the individuals seeking help.

A final note on the terminology used here; all of the dependent variables employed within this analysis are count data, which are of course non-negative. Generally, non-linear methods are avoided when performing causal inference within the DiD design since these estimates seldom represent the actual treatment effect. Unfortunately, OLS DiD estimates can be significantly biased when modelling non-negative variables (Lee & Lee, 2021),


and the magnitude of this bias can be substantial when outcomes are zero-inflated (as mortality and hospitalization data often is). To address this issue, I opt instead to employ a quasi-maximum likelihood Poisson estimator (QP) to compute my results. QP has a number of distinct advantages over OLS here: For one, estimating model 2.1 with QP-while nonlinear- still identifies the average treatment effect as β_{12} , albeit a multiplicative one. Lee and Lee (2021) outline the necessary assumptions for β_{12} to describe treatment effects, but they are analogous to a conventional DiD estimated with OLS, with the caveat that the actual causal estimator is really describing the ratio-in-ratios (RiR), not absolute differences. The RiR estimate therefore describes the causal elasticity (that is, the proportional relationship) between treatment and the outcome variable. When interpreting results I will be certain to emphasize this latter point, but for the remainder of this paper I will adopt language that is consistent with discipline-norms and use the term DiD flexibly to in fact mean RiR. Beyond providing me with a causal treatment measure, the semi-parametric nature of QP means it is capable of remaining a consistent estimator when data are zero-inflated (Lee & Lee, 2021); and simulations suggest that it outperforms mixed-models such as the zero-inflated Poisson or negative binomial regressions (Staub & Winkelmann, 2013) when the zero-inflation data generating process is unknown. Lastly, unlike most other nonlinear estimators, QP is generally not susceptible to the incidental parameters problem (Wooldridge, 1999; Weidner & Zylkin, 2021), so the two-way fixed effects estimator derived in equation 2.1 remains consistent.

2.4 Results

2.4.1 Primary results

To determine the causal impact of the East Liverpool misinformation incident on the actions of first responders, I utilize two distinct treatment specifications: The first contrasts naloxone administrations by LEOs after the East Liverpool event with those by COOPs during the same timeframe, while the second treats both EMS and LEO administrations in the post-event period as affected by misinformation. For the initial specification, the

Figure 2.2: Average naloxone administration counts timeseries by county-quarter and first responder type.



Average naloxone administrations, and naloxone administrations per opioid-related death by first responder type per county-quarter. Shaded regions represent the 95% confidence intervals and the dashed vertical line demonstrates the date of the East Liverpool event.

two-way fixed effects regression model described in equation (1) is calculated using a quasi-MLE Poisson count approach, with the findings presented in Table 2.2.

To summarize these results, I observe significant and substantial reductions in the number of naloxone administrations by LEOs after the East Liverpool incident. Across various model specifications, there is a consistent, statistically significant decline in naloxone administrations by LEOs when compared to COOP responders. The adjustment for time-varying county demographics and overall trends in opioid use doesn't significantly alter the coefficient estimates between specifications (1) and (2), suggesting that COOP administrations might be effectively accounting for these factors. Given that Poisson coefficients are interpreted as log-elasticities, the results from the preferred model specification (2) suggest that post-East Liverpool, LEOs were administering naloxone at a rate of approximately $e^{-0.1642} \approx 0.849$ times the previous rate. Notably, the size and significance of this effect remain robust even after introducing county-quarter and county-agent type fixed effects. In the context of specification (4), this translates to an 18% reduction in naloxone administrations by LEOs. These findings underscore that, even after adjusting for potential biases like variations in opioid demand, availability, or naloxone access within county-quarters, law enforcement's administration rates remain

significantly lower compared to COOPs.

Collectively, the results from Table 2.2 suggest that LEOs are significantly decreasing their naloxone administration activities after being exposed to the misinformation following the East Liverpool event. Although it cannot be definitively stated that this reduction is due to heightened reluctance stemming from the fentanyl scare, the timing of the changes seems to support this theory. In the following section, I will explore alternative explanations that might challenge this interpretation of the results. However, as it stands, these estimates offer convincing evidence that media events are influencing the behavior of first responders.

To test whether other first responders, besides just LEOs were affected by misinformation, I re-estimate all the models from Table 2.2 but consider both LEOs and EMS as treated. The results for these are listed in Table 2.3, but to summarize: Coefficient estimates remain negative and of roughly the same magnitude, but are no longer significant. From this, I conclude that EMS remained relatively unaffected by media coverage of the East Liverpool event when compared to LEOs, which is consistent with the make up of first responders that are most commonly profiled as victims of accidental fentanyl exposure in news items.

In Table 2.4 I perform a series of robustness checks by re-estimating all four of the primary model specifications under different conditions. To account for any heterogeneity in treatment effects between LEOs and EMS, I substitute an agent type indicator for the treatment definitions employed for the estimates outlined in Tables 2.2 and 2.3. This permits me to observe the comparative change in naloxone administration rates between LEOs and EMS against COOPs as the baseline. Additionally, I weight all observations by county-quarter population, to account for any further heterogeneity in data generating process. Panel (a) essentially replicates the estimates from Tables 2.2 and 2.3, and returns qualitatively similar results: LEOs are demonstrating large, statistically significant declines in naloxone administration, while EMS are appear to be unaffected. Panel (b) employs the ratio of naloxone administrations to opioid-related deaths to adjust for potential unseen

Table 2.2: Naloxone administration regression results, LEO as treated.

Dependent Variable: Model:	Unique Naloxone Administrations Count			
	(1)	(2)	(3)	(4)
Post \times Treat	-0.1642** (0.0745)	-0.1642** (0.0766)	-0.1642** (0.0744)	-0.1993** (0.0815)
Post	-0.4176*** (0.0909)	-0.5113*** (0.1279)		
Equifax		0.1578 (0.1064)		
% Medicaid Enroll		1.202 (1.264)		
Unemployment		0.1405** (0.0694)		
% Const. Employ.		-14.60 (18.04)		
Avg. Earnings		-0.0002 (0.0003)		
<i>Fixed-effects</i>				
First Responder Type	Yes	Yes	Yes	
Year	Yes	Yes		Yes
Quarter of Year	Yes	Yes		
County		Yes	Yes	
County \times Year		Yes		
County \times Period			Yes	Yes
First Responder Type \times County				Yes
<i>Varying Slopes</i>				
		Yes		
Observations	2,280	2,208	1,922	1,921
Squared Correlation	0.00416	0.61148	0.74240	0.92181

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

All estimates reflect results from the regression specified in equation (1) estimated as quasi-MLE Poisson count models. The treatment here is defined as LEO status following the East Liverpool event in Q2 of 2017. The dependent variable is the count of unique naloxone administrations performed by each first responder type (LEOs and COOP only) per county-quarter. Covariates include county-level quarterly estimates for the percent of population with a sub-prime credit score, percent of population eligible for Medicaid, unemployment rate, percent of laborforce employed in construction, and average household earning. Standard errors are clustered to the county-level, and in specification (2) I include both linear and quadratic county-level time trend parameters; all other fixed effects are detailed within the table.

Table 2.3: Naloxone administration regression results, LEO and EMS as treated.

Dependent Variable: Model:	Unique Naloxone Administrations Count			
	(1)	(2)	(3)	(4)
Post \times Treat	-0.1169 (0.0920)	-0.1169 (0.0937)	-0.1169 (0.0919)	-0.0953 (0.1190)
Post	-0.1185 (0.0940)	-0.1387 (0.1126)		
Equifax		0.0602 (0.0371)		
% Medicaid Enroll		1.104*** (0.3593)		
Unemployment		0.0460 (0.0336)		
% Const. Employ.		7.934 (10.31)		
Avg. Earnings		-2.12×10^{-5} (0.0001)		
<i>Fixed-effects</i>				
First Responder Type	Yes	Yes	Yes	
Year	Yes	Yes		Yes
Quarter of Year	Yes	Yes		
County		Yes	Yes	
County \times Year		Yes		
County \times Period			Yes	Yes
First Responder Type \times County				Yes
<i>Varying Slopes</i>				
		Yes		
Observations	3,420	3,372	3,321	3,317
Squared Correlation	0.16143	0.83634	0.85494	0.92950

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

All estimates reflect results from the regression specified in equation (1) estimated as quasi-MLE Poisson count models. The treatment here is defined as either LEO or EMS status following the East Liverpool event in Q2 of 2017. The dependent variable is the count of unique naloxone administrations performed by each first responder type per county-quarter. Covariates include county-level quarterly estimates for the percent of population with a sub-prime credit score, percent of population eligible for Medicaid, unemployment rate, percent of laborforce employed in construction, and average household earning. Standard errors are clustered to the county-level, and in specification (2) I include both linear and quadratic county-level time trend parameters; all other fixed effects are detailed within the table.

shifts in the effectiveness of resuscitations. The scenario might involve not just a reduction in the number of naloxone administrations but also a scenario where administrations occur yet are less timely and effective due to hesitancy. The outcomes in Panel (b) align with those in Panel (a) for LEOs, but now EMS shows significant reductions. This trend might suggest that EMS interventions are becoming less effective, although it could also be mirroring the overall downward trend in naloxone administrations by EMS as depicted in Figure 2.2. Lastly, in Panel (c) I replicate the protocol employed in estimating the results from Panel (a), but use a conventional OLS DiD estimator. Unsurprisingly, these results are very similar in nature, demonstrating an average decline of 3.8 fewer naloxone administrations per county-quarter performed by LEOs in the post-period.

2.4.2 Addressing threats to validity

Up to this point my analysis has operated under the assumption that COOP responders are less affected by misinformation shocks, and although I have provided evidence supporting this assumption, it's possible that this premise may not hold, potentially compromising the robustness of my findings. To refine the understanding of how misinformation about fentanyl exposure influences the behavior of all first responders, I adopt an additional difference-in-differences approach akin to the one used by Carrieri et al. (2019) in their study on the impact of a court ruling in Italy that falsely linked MMR vaccines to autism—a ruling widely disseminated on social media. Carrieri et al. (2019) aimed to assess the effect of this misinformation on vaccination rates, facing the challenge of identifying a control group in a context where social media's reach is pervasive. Their solution was to use regional variations in broadband internet access as an interaction term with their treatment variable, under the logic that individuals in less connected rural areas would be less influenced by social media. Since broadband availability is presumably not related to vaccine hesitancy other than through the channel of misinformation, this approach provides a means to isolate average treatment effects.

I build on this strategy and estimate a dosage-response model that uses the percent of households with internet access (drawn from the average 5-year, county-level ACS

Table 2.4: Naloxone administration robustness checks results

Model:	Dependent Variable:			
	(1)	(2)	(3)	(4)
<i>(a) Quasi-Poisson DiD, population-weighted</i>	Unique Naloxone Administrations Count			
Post × EMS	-0.0168 (0.1654)	-0.0168 (0.1684)	-0.0168 (0.1652)	0.0639 (0.2089)
Post × LEO	-0.2601*** (0.0368)	-0.2601*** (0.0375)	-0.2601*** (0.0367)	-0.2212** (0.0935)
Observations	3,420	3,372	3,321	3,317
Squared Correlation	0.15975	0.83454	0.85407	0.93073
<i>(b) Quasi-Poisson DiD, population-weighted</i>	Unique Naloxone Administrations per Opioid-Related Death			
Post × EMS	-0.3090** (0.1186)	-0.3090** (0.1208)	-0.3090** (0.1185)	-0.1756 (0.1417)
Post × LEO	-0.2252** (0.1045)	-0.2252** (0.1065)	-0.2252** (0.1044)	-0.1525 (0.1166)
Observations	3,420	3,324	2,769	2,757
Squared Correlation	0.20996	0.60808	0.82507	0.93736
<i>(c) OLS DiD, population-weighted</i>	Unique Naloxone Administrations Count			
Post × EMS	11.40 (9.550)	11.40 (9.721)	11.40 (9.539)	11.28 (9.539)
Post × LEO	-3.860** (1.640)	-3.860** (1.669)	-3.860** (1.638)	-3.881** (1.633)
Observations	3,420	3,420	3,420	3,420
R ²	0.36976	0.62226	0.64717	0.86017
Within R ²	0.00767	0.01373	0.01134	0.02771

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Results outlined above represent the heterogeneous treatment effects of the East Liverpool misinformation shock for EMS and LEOs using COOPs as the control. Each numbered column (1)-(4) corresponds to the model specifications estimated in Tables 2.2 and 2.3. Panel (a) estimates a weighted Quasi-Poisson DiD with the count of naloxone administrations as the dependent variable; Panel (b) estimates a weighted Quasi-Poisson DiD with the count of naloxone administrations per opioid-related death as the dependent variable; and Panel (c) estimates a weighted OLS DiD with the count of naloxone administrations as the dependent variable. All observations are weighted by the county-quarter population, and standard errors are clustered at the county-level.

estimates) as an exogenous source of variation and interact this with the POST and agent type indicators. By applying this model that is otherwise identical to my preferred specification, I can track the variations in naloxone administration across all first responder types as internet access—and consequently potential misinformation exposure—increases. Should the rise in possible misinformation exposure via social media lead to a heightened reluctance among first responders to offer assistance, one would expect to see significantly negative coefficients for these continuous treatment terms. The findings from this dosage-response analysis are presented in Panel (a) of Table 2.5. Both LEOs and EMS show statistically significant reductions in naloxone administrations during the post-misinformation period, suggesting a correlation between internet access and decreased intervention in overdoses. While COOP administrations do show a negative relationship with internet access enhancement, the effect is much smaller and not statistically significant at the 5% level. Collectively, these results seem to indicate that with rising potential exposure to misinformation, it is primarily LEOs and EMS who exhibit notable decreases in naloxone administrations. In appendix 2.8.2, I expand upon this model to delve deeper into the mechanism of online misinformation exposure. I do this by using time-varying Google search frequencies for terms associated with fentanyl misinformation panic as a substitute for the Post indicator to analyze the influence more closely.

Incorporating county-quarter fixed effects is intended to mitigate any county-level, time-specific variations, especially those related to the availability of naloxone. However, there's a concern that this method might not fully address potential discrepancies in how naloxone administrations are reported, which could be influenced by its increased accessibility. This situation could lead to a scenario where individuals are more discerning about when to engage emergency services for overdose incidents. If all first responder groups experienced a uniform decline in being contacted for assistance, the primary model should still reflect this adjustment. Conversely, if there's a notable disparity in how users seek assistance from LEOs versus COOPs following the East Liverpool misinformation shock, this behavioral change could be incorrectly attributed to LEOs ²⁶. The ideal approach would involve

²⁶What I am alternatively suggesting here is that the observed disparities in naloxone administration

Table 2.5: Alternative continuous treatment specifications to address threats to validity

Dependent Variable:	Unique Naloxone Administrations Count
<i>(a) Web-based misinformation exposure and reported administrations</i>	
Post \times % Internet Access \times LEO	-0.3872** (0.1577)
Post \times % Internet Access \times EMS	-0.2757*** (0.0704)
Post \times % Internet Access \times COOP	-0.1575 (0.1293)
Observations	3,372
Squared Correlation	0.83654
<i>(b) Community naloxone access and reported administrations</i>	
Pharmacies per 100k \times LEO	-0.1140 (218.8)
Pharmacies per 100k \times EMS	-0.0605 (218.8)
Pharmacies per 100k \times COOP	0.0425 (218.8)
Observations	2,688
Squared Correlation	0.86054

County-clustered standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Panel (a) leverages variation in internet access from the ACS five-year average to examine whether differential exposure to web-based misinformation across counties corresponds to a heightened reluctance to administer naloxone. By interacting the internet access measure with the Post and Agent Type indicators according to my preferred specification, this approach permits identification of heterogeneous treatment responses to misinformation shocks across all three agent types. Panel (b) utilizes registries of pharmacies distributing naloxone without prescription under the state standing order, the data for which contains the county-level number of pharmacies per one hundred thousand people annually from 2016 to 2019. I examine the influence of pharmacy involvement as a proxy for community naloxone access on first responder interventions by interacting these pharmacy counts with the responder type dummies, substituting the PostxTreat variable, to reassess my preferred model for the 2016-2019 subset.

analyzing naloxone administration based on responder type against the backdrop of its availability in the community, yet the absence of such detailed data poses a challenge.

To navigate this data limitation, I utilize archived registries of pharmacies participating in a naloxone distribution initiative under the state standing order that allows non-prescription access. This data, which provides the county-level number of pharmacies per one hundred thousand people annually from 2016 to 2019, offers a proxy for understanding naloxone distribution dynamics. In my analysis, I examine the influence of pharmacy involvement on first responder interventions by interacting these pharmacy counts with the responder type dummies, substituting the `PostxTreat` variable, to reassess my preferred model for the 2016-2019 subset. The outcomes of this regression analysis are presented in Panel (b) of Table 2.5. These coefficients can be understood as representing the variation in administration rates among various first responder groups in response to enhanced naloxone availability. Although there is an indication of reduced administrations associated with better pharmacy access for LEOs and EMS, these findings are not statistically significant at the 5% level. Similarly, the COOP indicator points to a slight uptick in administrations, yet this too lacks statistical significance. These outcomes, while not conclusive for establishing causal relationships, suggest that enhanced community access to naloxone doesn't markedly affect the differing administration rates among first responders.

An alternative hypothesis could be that users are not shifting towards unreported naloxone use among their peers but instead showing a preference for aid from COOPs over LEOs. In this scenario, the distribution of naloxone administrations would be skewed towards COOP interventions rather than those by LEOs, not due to censorship but due to a disproportionate reliance on COOPs. While this theory is intriguing, it seems implausible given that individuals requiring emergency assistance typically cannot choose their responder. Furthermore, this explanation fails to align with the simultaneous changes observed in naloxone administrations alongside the East Liverpool incident. To

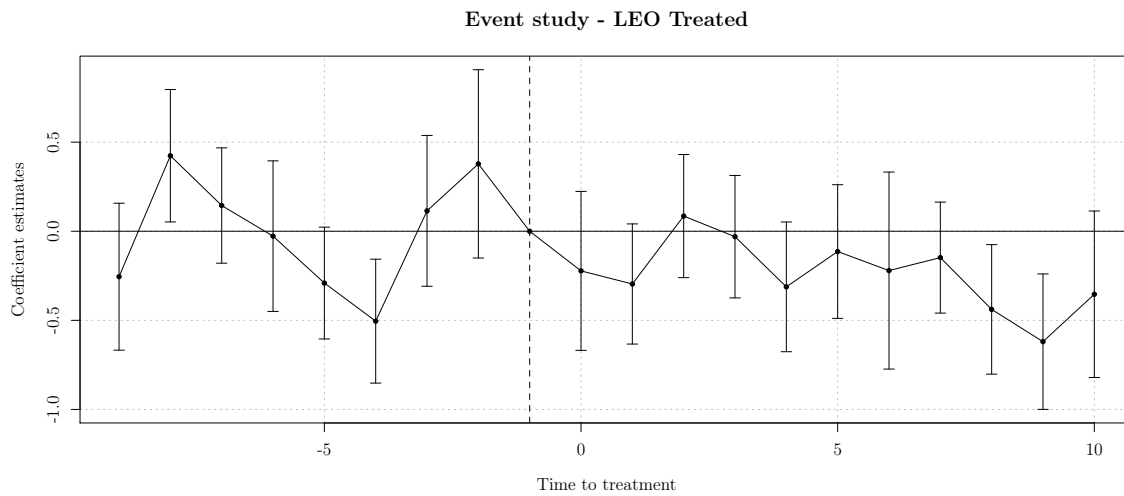
by different first responders are primarily due to a distinct reluctance to seek help from LEOs (perhaps out of concern of being arrested). This could be because, with the wider availability of naloxone kits, users may find less need to rely on LEOs for emergency assistance, a sentiment not mirrored in their interactions with COOPs or those trained by COOPs.

explore this theory further, I conduct an event study analysis applying my model to both scenarios where LEOs are treated and where both LEOs and EMS are treated, presenting the findings in Figure 2.3. Should there be noticeable pre-event trends that suggest a gradual shift from LEOs to COOPs prior to the misinformation incident, it would support this alternative explanation. However, the results predominantly show erratic pre-event patterns that stabilize into consistent declines following the misinformation event, particularly evident when examining combined treatment for LEOs and EMS, where significant drops in naloxone use emerge only post-East Liverpool.

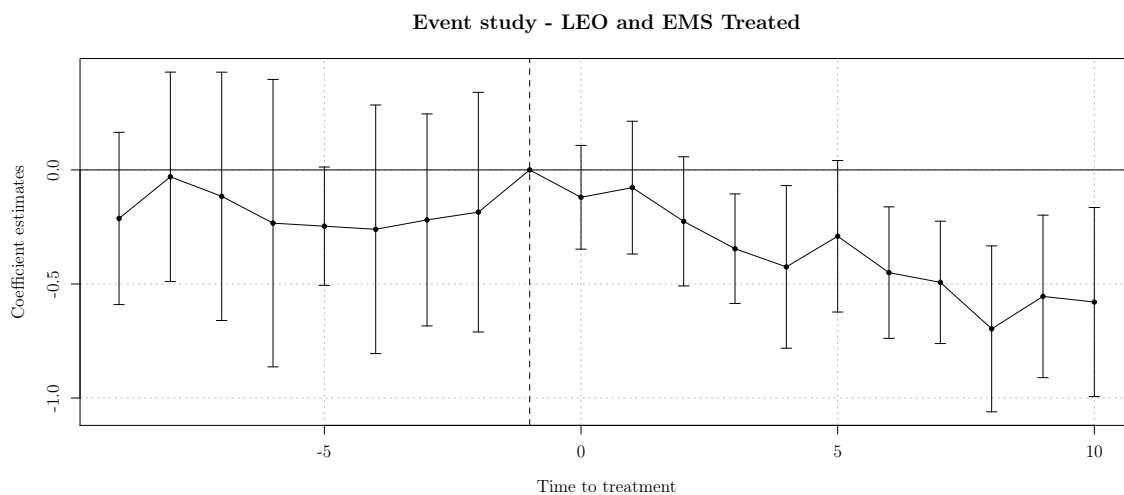
A similar concern with interpretation of my primary results is that COOP naloxone administration reporting likely declined in quality as kit distribution improved. Since laypersons cannot report directly to the state following a naloxone event, but only through their affiliated COOP, it's plausible that reporting compliance suffered as the number of trained individuals swelled. While this cannot be tested directly with the New York data employed here, surveys conducted with participants of other similar community-based opioid overdose prevention programs indicate a reluctance to notify EMS when naloxone is administered in private settings (Clark, Wilder, & Winstanley, 2014). Since the New York program specifically encourages opioid users to receive training through these COOPs, it may well be the case that these private-setting interventions are being largely unreported for similar reasons. Assuming no significant changes to LEO reporting compliance over the same time however, this actually suggests that these primary estimates are attenuated towards zero as it is underestimating the magnitude of interventions performed by laypersons.

2.5 Discussion

In the core part of my analysis, I have provided evidence supporting the idea that misinformation shocks related to fentanyl panic have led to a notable increase in reluctance among first responders to assist individuals experiencing an overdose, predominantly resulting in a reduced number of naloxone administrations by law enforcement officers.

Figure 2.3: Event studies for naloxone administrations, by treatment definition.

((a)) Event study results depicted here are according to the specification (4) from Table 2.2 and is estimating the relative decline in naloxone administrations performed by LEOs compared to COOP first responders. Models include county-quarter and county-first responder type fixed effects, standard errors are clustered at the county-level.



((b)) Event study results depicted here are according to the specification (4) from Table 2.3 and is estimating the relative decline in naloxone administrations performed by both EMS and LEOs compared to COOP first responders. Models include county-quarter and county-first responder type fixed effects, standard errors are clustered at the county-level.

While this information is crucial for understanding the underlying mechanism, it is equally vital to assess if these alterations in first responder actions have correspondingly resulted in higher rates of opioid-related deaths or hospitalizations. In an attempt to derive the overall effects of misinformation on mortality and hospitalizations, I estimate the following

regression:

$$Y_{c,t} = \beta \text{Post}_t + \Psi \text{Agent}_{c,t} + \Theta \text{Post}_t \times \text{Agent}_{c,t} + \delta X_{c,t} + \mu_c + \lambda_t + \varepsilon_{c,t}, \quad (2.2)$$

where $Y_{c,t}$ is the total number of opioid overdose deaths, ED visits, inpatient hospitalizations or rehab admissions in county c in quarter t ; $\text{Agent}_{c,t}$ is a vector composed of variables for the number of naloxone administrations performed by LEO, EMS and COOP, in that county-quarter; and $X_{c,t}$ is a vector of local socio-economic controls. Including independent counts of naloxone administrations performed by each first responder type, rather than aggregating to the county-quarter total permits varying efficacy. The parameter vector Ψ represents the elasticity estimates for naloxone administrations and the outcome variables. The Post indicator is intended to capture any post-treatment changes to mortality or hospitalization rates arising from misinformation exposure (perhaps as a result of less efficient or timely naloxone administration), and I will employ it flexibly when specifically attempting to identify such time variation. I would point out that, while the estimates derived in equation 2.1 are plausibly causal, there are likely other sources of endogeneity which prevent the parameters from vector Ψ from being viewed in the same way. Instead, I opt to treat these results as useful descriptive measures for the relationship between naloxone administration and mortality and hospitalization outcomes.

I estimate the regression outlined in equation 2.2 omitting the Post interaction term to derive the average effects of naloxone administration on hospitalization and rehab admissions across the entire 2015-2019 time span. Results are listed in Table 2.6. These models follow my preferred specification from Table 2.2 in column (2), with the notable omission of the quadratic county time-trend and first responder type fixed effect parameters, since I have only one observation for each county-quarter²⁷. In columns (1)-(3), I use ED visits by opioid type as the dependent variable; and in columns (4)-(6) I use (inpatient) hospitalizations. I observe slight decreases to ED visits associated with LEO naloxone

²⁷I would also note that, unlike opioid mortality and naloxone administrations, hospitalization and rehab admissions data is censored for county-quarter counts between 1-5, so these panels are both smaller and unbalanced compared to the naloxone administration estimates.

Table 2.6: Opioid hospitalization regression results

Dependent Variables:	ED Visits			Hospitalizations			Rehab Admissions	
	All Opioids (1)	Heroin (2)	Other Opioids (3)	All Opioids (4)	Heroin (5)	Other Opioids (6)	Heroin (7)	Other Opioids (8)
<i>Variables</i>								
LEO Naloxone	-0.0017** (0.0007)	-0.0021*** (0.0005)	-0.0018 (0.0017)	0.0016*** (0.0005)	0.0018** (0.0006)	0.0010* (0.0005)	-0.0003** (0.0001)	-0.0006* (0.0003)
EMS Naloxone	0.0034*** (0.0011)	0.0035*** (0.0012)	0.0026*** (0.0010)	0.0025** (0.0010)	0.0024 (0.0014)	0.0021** (0.0009)	0.0004** (0.0002)	-0.0004 (0.0003)
COOP Naloxone	0.0016** (0.0007)	0.0021*** (0.0008)	0.0010 (0.0009)	0.0006 (0.0012)	-0.0021 (0.0024)	0.0015* (0.0008)	-9.83×10^{-6} (0.0002)	-0.0004 (0.0003)
<i>Fit statistics</i>								
Observations	758	604	324	346	158	234	935	935
Squared Correlation	0.97540	0.96564	0.96736	0.96435	0.89889	0.94983	0.99816	0.99467

Clustered (county) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Results outlined above represent the heterogeneous elasticities of opioid-related ED (emergency department) visits and hospitalization to EMS, LEO, COOP naloxone administrations. These models are estimated according to equation (2) without the Post indicator and include county-level quarterly estimates for the percent of population with a sub-prime credit score, percent of population eligible for Medicaid, unemployment rate, percent of laborforce employed in construction, and average household earning. Fixed effects include county, county-quarter of year (season), and county-year, as well as a unique linear time trend for each county; standard errors are clustered at the county-level.

administrations, while the opposite is true for both EMS and COOPs. In terms of magnitude, this corresponds to a 0.17% decline in ED visits for all opioids per LEO naloxone administration, and 0.34% or 0.16% increases in ED visits per EMS and COOP naloxone administration. Hospitalizations are positively associated with increased naloxone administration among LEOs and EMS, with respective approximate declines of 0.16% and 0.25% per unique naloxone administration performed. Given that visits to the emergency department (ED) and hospital admissions correspond to more severe instances of overdose, it is plausible to assume that without the administration of naloxone, such cases could have been fatal. Therefore, the notably positive coefficients observed might be tentatively viewed as an implicit reduction in death rates.

In columns (7)-(9) I estimate the effects of naloxone administrations on unique rehab clinic admissions by substance dependency type and observe similarly significant, albeit modest declines resulting from LEO interventions. These highlight the complexity of such situations: While it may be socially-preferable to have overdosing individuals resuscitated to dying, and regular opioid users voluntarily seek rehab services to not, LEO interventions

that save lives also appear to reduce the likelihood of voluntary rehab admission. While one cannot be certain without precise data on post-intervention outcomes, it seems the most likely explanation here is the tendency for revived individuals to be incarcerated when they opt not to receive medical care. Surveys of post-naloxone administration LEO protocols corroborate this theory (Smiley-McDonald et al., 2022).

I also estimate a similar model examining opioid-related mortality and report the results in Table 2.7. To summarize: I use raw counts of opioid mortality, by opioid type as the dependent variable and find null to slightly *positive* associations between naloxone administration and mortality. In terms of magnitude, these EMS and COOP estimates vary from 0.3% to 0.4% increases to mortality per naloxone administration. The coefficient estimates appear to suggest that the number of opioid-related deaths increases with more naloxone administrations. Clearly, the number of administrations is endogenous with overall mortality, as an increase in total overdoses is correlated with the number of doses administered. In appendix 2.8.1 I explore this more by re-estimating model (1), but with the dependent variable as the ratio of opioid-related deaths to opioid-related hospitalizations. To summarize those results, I observe similar positive increases to mortality when accounting for ED visits and hospitalizations, though the statistically significant coefficients are predominately concentrated among LEOs.

Determining the economic magnitude of these effects on mortality is somewhat problematic. As I outlined above, the mortality estimates derived from my secondary estimates are probably biased to the point of flipping signs, and as such not reliable. Additionally, while significant, the hospitalization results are suspiciously small in magnitude. Nonetheless, doing some back-of-the-envelope calculations with an average number of 5.76 LEO naloxone administrations per county in the post-period, then the 15.1% decline in administration rate implies 0.872 fewer administrations per county-quarter. Using the population-weighted model results from Panel (a) of Table 2.4, these estimates imply a 22.9% decline in naloxone interventions performed by LEOs, or 8.74 fewer unique administrations per county in the post-period. Using the secondary heroin estimates, this would imply approximately 1 less ED visit, and- if hospitalizations are interpreted as avoided deaths- 1 less death

Table 2.7: Opioid mortality regression results

Dependent Variables: Model:	Overdose Deaths		
	All Opioids (1)	Heroin (2)	Other Opioids (3)
LEO Naloxone	0.0007 (0.0011)	0.0019 (0.0012)	0.0003 (0.0013)
EMS Naloxone	0.0034* (0.0019)	0.0032** (0.0014)	0.0037* (0.0021)
COOP Naloxone	0.0032*** (0.0008)	0.0043** (0.0020)	0.0040*** (0.0008)
<i>Fit statistics</i>			
Observations	1,096	883	1,081
Squared Correlation	0.95942	0.91893	0.94909

Clustered (county) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Results outlined above represent the heterogeneous elasticities of opioid-related mortality, by opioid type, to EMS, LEO, COOP naloxone administrations. These models are estimated according to equation (2) without the Post indicator and include county-level quarterly estimates for the percent of population with a sub-prime credit score, percent of population eligible for Medicaid, unemployment rate, percent of laborforce employed in construction, and average household earning. Fixed effects include county, county-quarter of year (season), and county-year, as well as a unique linear time trend for each county; standard errors are clustered at the county-level.

across the entire state in the post-period. These estimates seem far too low, considering anecdotal success with which LEO naloxone administrations appear to be in saving lives (Smiley-McDonald et al., 2022). For instance, Pourtaher et al. (2022) found that in New York state, 87.4% of overdose victims were saved in incidents where LEOs intervened; and among the victims who did not survive, only 5% were still alive upon police arrival, before EMS reached the scene. If each county averaged over eight fewer naloxone interventions in the post-period, a magnitude estimate between 1-4 fewer deaths *per county* seems more feasible. Clearly, without a means of comprehensively addressing endogeneity sources in the mortality models, these cumulative effects are going to be biased downwards.

Despite the absence of direct proof connecting the decrease in naloxone administration to an increase in opioid overdose deaths or hospitalizations, the results underscore an important policy takeaway from this analysis. They point out the potential flaws in previous studies that investigated how naloxone access laws (NALs) influence economic outcomes, attempting to determine naloxone's causal effect on mortality and morbidity. During the period between 2015 and 2019 covered by this research, twenty-three states enacted their inaugural NALs (Prescription Drug Abuse Policy System, 2023). Should the behavioral reactions of first responders in these states mirror those observed in New York following the 2017 misinformation event, it could introduce a major variable that complicates the interpretation of NALs' effectiveness. If the observed reluctance of first responders to use naloxone coincides with the implementation of these harm reduction measures, then any lack of observable results could mistakenly be attributed to the unintended consequences of moral hazard. To illustrate the scale, Doleac and Mukherjee (2018) noted a 15% rise in opioid-related emergency department (ED) visits in states after implementing naloxone access laws (NALs). If the findings from New York can be extrapolated to the entire U.S., then the observed 15.1% to 22.9% reduction in naloxone administrations by LEOs—coupled with the fact that LEOs are typically the first responders to overdose incidents—could significantly challenge the moral hazard explanation for the surge in ED visits. Indeed, Kochersperger (2023) highlights a similar observation, showing varied opioid-related mortality rates in reaction to the media coverage surrounding the East

Liverpool event, which suggests that this explanation holds comparable plausibility.


This observation not only points out a potential flaw in such policies but also suggests different directions for remedial actions: Should misinformation about the risks of fentanyl significantly elevate mortality via the behavioral responses of first responders, emphasis should be placed on debunking these myths and enhancing training for first responders. Conversely, if the easier access to naloxone is leading to riskier behaviors among users, this would raise questions about the overall effectiveness of harm reduction strategies writ large. In an ideal scenario, this research here would leverage the design examining first responder misinformation susceptibility to evaluate the competing theories, but the absence of variation in NAL policy during the relevant data period precludes this approach. Consequently, this limitation underscores the necessity for more nuanced investigations into NAL effectiveness, focusing on concrete indicators of naloxone distribution and utilization, rather than solely on policy changes.

Another interesting research question for further analysis is to determine what sort of factors are influencing treatment heterogeneity. In Figure 2.4 I re-estimate model (4) from my primary naloxone administration results in Table 2.3, but interact the treatment indicator with a vector of county indicators to back-out average, county-level treatment effects; I then compute the percent change in the *total* number of naloxone administrations (that is, including LEO, EMS and COOPs) observed relative to the counterfactual²⁸. While most counties did experience significant decreases in naloxone administration rates overall, a few actually appear to have increased over counterfactual scenario. At first thought, this may have something to do with the degree of urbanization, but visual inspection reveals no obvious trend. Doleac and Mukherjee (2018) describe variation between urban and rural areas in opioid arrests following adoption of NALs, but are not able to observe such outcomes for New York or the Northeastern United States more generally. In line with their ambiguous results, I will concede that I cannot say with certainty whether this is a significant factor here. Instead of interrogating this matter

²⁸I drop a handful of counties which recorded no COOP administrations in the post-period; they are shaded grey

further, I leave this curiosity as an open question now for future research.

Figure 2.4: County-Level Heterogenous Treatment Effects.



Color scale represents the percent change in the observed number of naloxone administrations compared to the counterfactual scenario. Shaded grey areas represent urban areas, while solid-grey counties area omitted due to lack of post-treatment data.

2.6 Conclusions

This research has provided a detailed exploration of the interplay between opioid-related mortality and the behaviors of first responders in the context of the opioid crisis, with a particular focus on the effect of misinformation on naloxone administration. Employing a robust dataset from the New York State Department of Health, which segregates naloxone administration instances by type of first responder, I am able to isolate the impact of a significant public health misinformation event in 2017. This approach is able to identify the causal effect of misinformation on the likelihood of naloxone use by different first responders, particularly contrasting law enforcement officers with harm reduction professionals. The results reveal a discernible decrease in naloxone administrations by police officers post-misinformation, suggesting that erroneous beliefs about the dangers of fentanyl exposure significantly deter its use, despite the availability and known efficacy of

naloxone in reversing opioid overdoses.

The implications of this research, especially when taken together with observations from Kochersperger (2023), touch upon the very strategies employed to combat the opioid crisis. By revealing the substantial impact of misinformation on the actions of first responders, particularly law enforcement officers, this study underscores a critical barrier to the effective implementation of naloxone access laws (NALs). While these laws are designed to enhance naloxone accessibility and encourage its use, the entrenched misconceptions regarding the risks associated with opioid exposure, especially fentanyl, can stifle these efforts. This hesitancy not only affects the immediate response to overdose incidents but also has broader implications for public health strategies aimed at reducing opioid-related fatalities. Therefore, addressing misinformation and improving education among first responders emerge as essential components in maximizing the effectiveness of harm reduction policies.

Moreover, this analysis adds a new layer to the discourse on the efficacy of NALs, suggesting that the obstacles to their success are not merely legal or logistical but also psychological and informational. The study's findings advocate for a more holistic approach to harm reduction that incorporates not just legal reforms and medical interventions but also targeted information campaigns to dispel myths and educate first responders about the realities of opioid exposure risks. By doing so, policymakers can ensure that naloxone's life-saving potential is fully realized, contributing to a more informed and effective response to the opioid epidemic. Experimentation with such retraining efforts have been broadly successful (Winograd, Phillips, et al., 2020; Del Pozo et al., 2021), but the existence of some mixed results hints at the need for further study (Winograd, Stringfellow, Phillips, & Wood, 2020). This approach could bridge the gap between policy intent and on-the-ground impact, offering a more nuanced understanding of the dynamics at play in the fight against this pervasive public health challenge.

References

- Adams, A., Maloy, C., & Warrick, B. J. (2023). Does occupational exposure to fentanyl cause illness? a systematic review. *Clinical Toxicology*, *61*(9), 631–638.
- Attaway, P. R., Smiley-McDonald, H. M., Davidson, P. J., & Kral, A. H. (2021). Perceived occupational risk of fentanyl exposure among law enforcement. *International Journal of Drug Policy*, *95*, 103303.
- Beletsky, L., Seymour, S., Kang, S., Siegel, Z., Sinha, M. S., Marino, R., . . . Freifeld, C. (2020). Fentanyl panic goes viral: The spread of misinformation about overdose risk from casual contact with fentanyl in mainstream and social media. *International Journal of Drug Policy*, *86*, 102951.
- Berardi, L., Bucerius, S., Haggerty, K. D., & Krahn, H. (2021). Narcan and narcan't: Implementation factors influencing police officer use of narcan. *Social science & medicine*, *270*, 113669.
- Bucerius, S., Berardi, L., Haggerty, K. D., & Krahn, H. (2022). New drugs, new fears: synthetic opioids and adaptations to police practice. *Policing and Society*, *32*(7), 832–845.
- Buresh, M., Genberg, B. L., Astemborski, J., Kirk, G. D., & Mehta, S. H. (2019). Recent fentanyl use among people who inject drugs: results from a rapid assessment in baltimore, maryland. *International Journal of Drug Policy*, *74*, 41–46.
- Carrieri, V., Madio, L., & Principe, F. (2019). Vaccine hesitancy and (fake) news: Quasi-experimental evidence from italy. *Health economics*, *28*(11), 1377–1382.
- Clark, A. K., Wilder, C. M., & Winstanley, E. L. (2014). A systematic review of community opioid overdose prevention and naloxone distribution programs. *Journal of addiction medicine*, *8*(3), 153–163.
- Del Pozo, B., Sights, E., Kang, S., Goulka, J., Ray, B., & Beletsky, L. A. (2021). Can touch this: training to correct police officer beliefs about overdose from incidental contact with fentanyl. *Health & Justice*, *9*(1), 1–6.
- Doleac, J. L., & Mukherjee, A. (2018). The moral hazard of lifesaving innovations:

- naloxone access, opioid abuse, and crime.
- Erfanian, E., Grossman, D., & Collins, A. R. (2019). The impact of naloxone access laws on opioid overdose deaths in the us. *Review of Regional Studies*, *49*(1), 45–72.
- Hedegaard, H., Miniño, A. M., Spencer, M. R., & Warner, M. (2021). Drug overdose deaths in the united states, 1999–2020.
- Herman, P. A., Brenner, D. S., Dandorf, S., Kemp, S., Kroll, B., Trebach, J., . . . Stolbach, A. I. (2020). Media reports of unintentional opioid exposure of public safety first responders in north america. *Journal of Medical Toxicology*, *16*(2), 112–115.
- Kochersperger, E. (2023). Quantifying the Effects of Fentanyl Exposure Misinformation on Opioid Mortality. *Working Paper*.
- Kucher, K., & Figueroa, T. (2021, Aug). *Cops were taught that incidental fentanyl exposure is deadly. experts say it's unlikely. why the disconnect?* Retrieved from <https://www.sandiegouniontribune.com/news/public-safety/story/2021-08-29/cops-were-taught-that-incidental-fentanyl-exposure-is-deadly-experts-say-its-unlikely-why-the-disconnect>
- Lee, M.-j., & Lee, S. (2021). Difference in differences and ratio in ratios for limited dependent variables. *arXiv preprint arXiv:2111.12948*.
- Lynch, M. J., Suyama, J., & Guyette, F. X. (2018). Scene safety and force protection in the era of ultra-potent opioids. *Prehospital emergency care*, *22*(2), 157–162.
- McBride, A. J., Pates, R. M., Arnold, K., & Ball, N. (2001). Needle fixation, the drug user's perspective: a qualitative study. *Addiction*, *96*(7), 1049–1058.
- McClellan, C., Lambdin, B. H., Ali, M. M., Mutter, R., Davis, C. S., Wheeler, E., . . . Kral, A. H. (2018). Opioid-overdose laws association with opioid use and overdose mortality. *Addictive behaviors*, *86*, 90–95.
- Moss, M. J., Warrick, B. J., Nelson, L. S., McKay, C. A., Dubé, P.-A., Gosselin, S., . . . Stolbach, A. I. (2018). Acmt and aact position statement: preventing occupational fentanyl and fentanyl analog exposure to emergency responders. *Clinical Toxicology*, *56*(4), 297–300.
- NYSDOH. (2023). *New york state opioid annual data report*. Retrieved

- from https://www.health.ny.gov/statistics/opioid/data/pdf/nys_opioid_annual_report_2023.pdf
- Persaud, E., & Jennings, C. R. (2020). Pilot study on risk perceptions and knowledge of fentanyl exposure among new york state first responders. *Disaster medicine and public health preparedness*, *14*(4), 437–441.
- Pourtaher, E., Payne, E. R., Fera, N., Rowe, K., Leung, S.-Y. J., Stancliff, S., . . . Dailey, M. W. (2022). Naloxone administration by law enforcement officers in new york state (2015–2020). *Harm reduction journal*, *19*(1), 102.
- Prescription Drug Abuse Policy System. (2023). *Naloxone Overdose Prevention Laws*. <https://pdaps.org/datasets/opioid-analgesics-prescribing-limits>.
- Rees, D. I., Sabia, J. J., Argys, L. M., Dave, D., & Latshaw, J. (2019). With a little help from my friends: the effects of good samaritan and naloxone access laws on opioid-related deaths. *The Journal of Law and Economics*, *62*(1), 1–27.
- Rouhani, S., Park, J. N., Morales, K. B., Green, T. C., & Sherman, S. G. (2019). Harm reduction measures employed by people using opioids with suspected fentanyl exposure in boston, baltimore, and providence. *Harm reduction journal*, *16*(1), 1–9.
- Smart, R., Pardo, B., & Davis, C. S. (2021). Systematic review of the emerging literature on the effectiveness of naloxone access laws in the united states. *Addiction*, *116*(1), 6–17.
- Smiley-McDonald, H. M., Attaway, P. R., Richardson, N. J., Davidson, P. J., & Kral, A. H. (2022). Perspectives from law enforcement officers who respond to overdose calls for service and administer naloxone. *Health & Justice*, *10*(1), 9.
- Staub, K. E., & Winkelmann, R. (2013). Consistent estimation of zero-inflated count models. *Health economics*, *22*(6), 673–686.
- Tompkins, C., Ghoneim, S., Wright, N., Sheard, L., & Jones, L. (2007). Needle fear among women injecting drug users: a qualitative study. *Journal of Substance Use*, *12*(4), 281–291.
- Weidner, M., & Zylkin, T. (2021). Bias and consistency in three-way gravity models. *Journal of International Economics*, *132*, 103513.

- White, M. D., Watts, S., Orosco, C., Perrone, D., & Malm, A. (2022). Leveraging body-worn camera footage to better understand opioid overdoses and the impact of police-administered naloxone. *American journal of public health, 112*(9), 1326–1332.
- Winograd, R. P., Phillips, S., Wood, C. A., Green, L., Costerison, B., Goulka, J., & Beletsky, L. (2020). Training to reduce emergency responders' perceived overdose risk from contact with fentanyl: early evidence of success. *Harm Reduction Journal, 17*(1), 1–5.
- Winograd, R. P., Stringfellow, E. J., Phillips, S. K., & Wood, C. A. (2020). Some law enforcement officers' negative attitudes toward overdose victims are exacerbated following overdose education training. *The American journal of drug and alcohol abuse, 46*(5), 577–588.
- Wooldridge, J. M. (1999). Quasi-likelihood methods for count data. *Handbook of applied econometrics volume 2: Microeconomics, 321–368.*

2.8 Appendices

2.8.1 Further mortality analysis, controlling for overdose counts

I estimate the regression outlined in equation 2.2 omitting the Post interaction term to derive the average effects of naloxone administration on opioid-related deaths per ED visits and hospitalizations. These models follow my preferred specification from Table 2.2 in column (2), with the notable omission of the quadratic county time-trend and first responder type fixed effect parameters. The results are listed in Table 2.8.

Column (1) is identical to that of Table 2.7; while columns (2) and (5) use the ratio of deaths to ED visits as the dependent variable; and columns (3) and (6) use the ratio of deaths to total hospitalizations (that is, the sum of opioid overdose ED visits and hospitalizations). For the model results depicted in columns (4)-(5), I additionally include a Post interaction term for LEOs and EMS to capture any changes in naloxone administration efficacy that extend beyond the extensive margin of dose counts (arising, say, from the delay of naloxone administration, rather than outright refusal). By omitting a Post \times COOP interaction, these coefficients should be able to identify the change in relative elasticity of opioid-related mortality to naloxone administration following the East Liverpool misinformation shock using COOPs as a baseline. This is essentially the same identification assumption employed in the primary naloxone administration analysis, but here I assume that in the counterfactual scenario LEO and EMS naloxone administrations should change at a proportionally identical rate to COOPs in the post-period. Observe that COOP and EMS coefficients are generally slightly negative as expected, though insignificant. LEO coefficients are not significantly positive, and even slightly larger in magnitude.

There could be several reasons for these unanticipated mortality results. It would appear that LEO administrations have been more effective at preventing deaths following treatment, perhaps as a result of the misinformation treatment itself. More plausibly however, it could be that- as observed from the naloxone administration regression results-

Table 2.8: Mortality results, controlling for overdose counts.

Dependent Variables:	Deaths	Deaths per ED Visit	Deaths per Tot. Hosp.	Deaths	Deaths per ED Visits	Deaths per Tot. Hosp.
Model:	(1)	(2)	(3)	(4)	(5)	(6)
<i>Variables</i>						
LEO Naloxone	0.0007 (0.0011)	0.0040*** (0.0013)	0.0026** (0.0010)	0.0024** (0.0010)	0.0081*** (0.0029)	0.0058*** (0.0018)
EMS Naloxone	0.0034* (0.0019)	-0.0015 (0.0013)	-0.0015 (0.0013)	0.0025 (0.0020)	-0.0024 (0.0017)	-0.0026* (0.0014)
COOP Naloxone	0.0032*** (0.0008)	5.88×10^{-5} (0.0017)	-0.0009 (0.0017)	0.0024*** (0.0008)	-0.0008 (0.0016)	-0.0015 (0.0015)
Post				-0.1483 (0.1352)	-0.0484 (0.1671)	-0.1346 (0.1330)
LEO \times Post				-0.0034* (0.0019)	-0.0089** (0.0041)	-0.0069*** (0.0025)
EMS \times Post				0.0012 (0.0010)	0.0006 (0.0017)	0.0015 (0.0012)
<i>Fit statistics</i>						
Observations	1,096	729	366	1,096	729	366
Squared Correlation	0.95942	0.74291	0.83784	0.96101	0.74820	0.84247

Clustered (county) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Results outlined above represent the heterogeneous elasticities of total opioid-related mortality per ED (emergency department) visit and hospitalizations to EMS, LEO, COOP naloxone administrations. These models are estimated according to equation (2) both with and without the Post indicator and include county-level quarterly estimates for the percent of population with a sub-prime credit score, percent of population eligible for Medicaid, unemployment rate, percent of laborforce employed in construction, and average household earning. The Post interaction term for LEOs and EMS is meant to capture any changes in naloxone administration efficacy from the delay of naloxone administration, rather than outright refusal. Fixed effects include county, county-quarter of year (season), and county-year, as well as a unique linear time trend for each county; standard errors are clustered at the county-level.

LEOs are less inclined to administer naloxone, and that some degree of selection is occurring which is pushing those first responders which would otherwise delay administration towards not administering at all. Since misinformation appears to be most influential in the absence of accurate information on opioid exposure hazards, one could describe individual treatment heterogeneity as arising from some latent variable for general opioid literacy. In such a conceptual setting, it is reasonable to presume that those LEOs which would be the least affected by misinformation are also more capable of employing naloxone more effectively.


Without a more complete measure for the total number of potentially-avertable lethal overdoses that LEOs respond to, this positive coefficient is probably biased upward as it captures exogenous changes to opioid lethality and the general excludability by agent-type of naloxone administration (that is, naloxone administration counts are plausibly correlated across agent-types). With regards to that latter point, consider for instance that if an LEO administers naloxone when EMS are present, this would effectively increase LEO counts while lowering EMS'. Since LEOs are consistently arriving to overdose scenes sooner than EMS or COOPs (Pourtaher et al., 2022), then it may be that a larger portion of those resuscitation attempts are futile. As well, it may be that EMS are better at establishing whether an overdose victim is past the point of successful resuscitation, and opt to use naloxone less often in such scenarios. Lastly, it is simply difficult to envision a scenario where an increase in naloxone administration is *causing* an increase in opioid-related mortality. There are no doubt moral hazard arguments that the increased availability of naloxone may induce more reckless behavior among regular opioid users (Doleac & Mukherjee, 2018), but it is unclear why such effects would only manifest in relation to LEO administrations as they do here.

2.8.2 Additional web-based exposure results


I replicate the web access dosage-response protocol outlined in Panel (a) of Table 2.5, but additionally relax the treatment periods to represent the comparative frequency of Google searches for fentanyl exposure hazards. Figure 2.6 represents data taken from Kochersperger (2023) and shows a clear change in search patterns following the East

Liverpool event, but these treatment dynamics may still be more complex than our on-off design is capable of describing. For instance, if some proportion of these Google searches correspond to one unique individual becoming aware of this misinformation, then the *cumulative* index up to some point in time should more accurately capture this dissemination of information. This is probably more reasonable (this dissemination process to take time, after all), but adopting such an approach allows me to exploit a greater source of temporal variation as well. As well, this approach allows me to account for the cyclical nature of social media coverage related to fentanyl misinformation panic events, whereby individuals periodically re-share older news articles as new events occur.

I estimate a dosage-response model that uses the interaction of percent of households with internet access (drawn from the average 5-year, county-level ACS estimates), Google search trends, and the LEO treatment indicator from equation 2.1 to fully utilize both spatial and temporal variation in treatment exposure. I find that this dosage-response coefficient is highly significant and negative ($\beta_D = -0.006015$, standard error = 0.00125), suggesting the increased likelihood of having queried Google for information pertaining to fentanyl exposure myths within a given county-quarter corresponds to a marginal decrease in the number naloxone administrations performed by LEO, relative to COOPs. Direct, quantitative interpretation of this coefficient is difficult, however, so instead I compute the counterfactual scenario where this exposure measure has a null effect and plot these values against the observed LEO administrations in Figure 2.6. While I do not formally employ the date of the East Liverpool event as a treatment, observe that the observed and counterfactual naloxone administrations remain very close to one another until just after that occurred, lending credence to my primary identification strategy.

Figure 2.5: Time series of Google search interest in the hazards of fentanyl exposure.

Google search data was collected from Google Trends and represents the relative popularity of search terms over the specified time frame; the time series was derived from querying the Google trends for “fentanyl AND (touch* OR contact* OR absor* OR inhal* OR expos*)”. Solid black line is the weekly average for the Google Trends Index, while the dashed blue line is the rolling average of the 20 preceding weeks. The dashed vertical line demonstrate the date of the East Liverpool event.

Figure 2.6: Average LEO naloxone administrations, observed and counterfactual.

Counterfactual values are computed by dividing the observed number of LEO naloxone doses in county c in quarter t by $e^{\beta_D \times IA_c \times GT_t}$, where IA_c is the percent of households in county c with regular internet access; GT_t is the Google search trends index for the relative search prevalence for fentanyl misinformation in period t ; and β_D is the dosage-response coefficient derived from estimating equation 2.1 after interacting the LEO treatment indicator with both percent internet access and Google search trends. Shaded regions represent the 95% confidence intervals.

3 Insurance Barrier Impacts on Vaccine Hesitancy: Administrative Burden and COVID-19 Vaccination within the Medicaid Population

3.1 Introduction

The lower engagement in preventative care among Medicaid recipients, compared to those with private insurance, signals a nuanced challenge in public health accessibility and equity. Medicaid expansions, while instrumental in extending healthcare coverage, have not fully bridged the gap in preventative care utilization (H. Allen, Gordon, Lee, Bhanja, & Sommers, 2021). This persistent shortfall raises concerns about the long-term health implications for Medicaid beneficiaries, who may face elevated risks of chronic conditions and delayed diagnosis of serious illnesses. As this issue unfolds within the broader healthcare landscape, understanding the underlying factors contributing to this gap is crucial for crafting effective interventions that can elevate preventative care rates to those comparable with privately insured populations.

The intricate web of administrative procedures associated with Medicaid enrollment and retention stands out as a significant obstacle to accessing preventative care. For many, the daunting task of navigating these processes can deter engagement with essential healthcare services, leading to underutilization (Moynihan, Herd, & Ribgy, 2016). Beyond the bureaucratic complexities, other factors such as provider availability, awareness of eligible services, and social determinants of health play pivotal roles in shaping healthcare behaviors among Medicaid recipients (E. M. Allen, Call, Beebe, McAlpine, & Johnson, 2017). Delving into these aspects, it becomes evident that the issue is multifaceted, involving not just systemic and procedural barriers, but also personal and community-level factors that collectively influence healthcare utilization patterns.


This study underscores the importance of understanding the diverse factors that impact preventative care utilization among Medicaid recipients, with special emphasis placed

on the role of administrative burden in this process. Utilizing the COVID-19 pandemic as a case study, the central research question explores the extent to which Medicaid enrollment status influences COVID-19 vaccination decisions. It delves into whether low vaccination rates observed among Medicaid recipients²⁹ stem from an overrepresentation of vaccine-hesitant populations within Medicaid or if enrollment in the program itself alters vaccination behaviors. If there is a direct impact, the study seeks to understand the direction of this influence and the driving factors behind it. A key aspect of this inquiry is whether vaccine reluctance in the Medicaid cohort is rooted in skepticism about the costs or benefits of vaccination, shaped by their experiences with Medicaid enrollment and care. Additionally, this research aims to identify the most suitable comparison group for assessing Medicaid outcomes—exploring whether policymakers should benchmark program effectiveness against the uninsured or the insured to gauge success or failure.

To perform this analysis I will employ the U.S. Census Bureau’s Household Pulse Survey to observe differences in the marginal likelihoods to vaccinate between Medicaid recipients and demographically-similar uninsured and privately insured populations. To mitigate potential selection biases, such as the inclination for Medicaid recipients to possess poorer health status or belong to racial and ethnic groups with lower trust in government policies, I utilize a variety of empirical methods. These approaches aim to generate a more representative counterfactual, taking into account the complex dynamics within the Medicaid population. My findings generally support the assertion that an increase in administrative burden associated with managing state Medicaid programs correlates with a reduced likelihood of vaccination against COVID-19 among its recipients. This pattern emerges clearly when contrasting Medicaid recipients with those covered by other insurance providers; however, Medicaid recipients are more inclined to get vaccinated compared to the uninsured, underscoring the program’s benefits. A deeper exploration into the specific reasons for vaccine hesitancy among Medicaid recipients indicates a lower trust in the COVID-19 vaccines. Overall, the results suggest that while the Medicaid population may have more familiarity with navigating the healthcare system, they are

²⁹See: <https://kffhealthnews.org/news/article/medicaid-covid-vaccine-obstacles-states/>

Figure 3.1: County-level primary vaccination series completion rate time series, stratified by Medicaid enrollment rates.




Plotted lines correspond to the percent of county populations recorded as having completed the primary series (initial and first booster dose) of COVID-19 vaccinations over time, with the red and blue lines corresponding to counties below the 25th, or above the 75th percentiles for percentages of population enrolled in Medicaid. Vaccination data is drawn from the CDC’s compilation of immunization information systems reported vaccinations. Shaded regions correspond to the 95% confidence bounds. To address qualitative differences in population characteristics and demographics between high- and low-Medicaid enrollment counties, the black dashed line represents the entropy balanced average vaccination rate among the 25th percentile counties that best approximates the 75th percentile counties. Balanced covariates include county population, median age and household income, percents of population white, black, and having completed college; all balancing data are drawn from the 2021 5-year ACS estimates.

less likely to get vaccinated compared to those with private insurance, potentially due to the anticipated time and effort involved in the vaccination process.

The motivation for this avenue of research is drawn from the growing literature that has examined the role that administrative burden (Stuber & Kronebusch, 2004; E. M. Allen et al., 2017; Fox, Stazyk, & Feng, 2020) and cost-sharing (Wright et al., 2005; E. M. Allen et al., 2017) has played on discouraging Medicaid enrollments and care utilization. Due to the bureaucratic burdens of enrolling, social stigma, unanticipated cost-sharing, difficulties in accessing care, and discrimination by providers, Medicaid recipients may naturally become reluctant to utilize care services, such as vaccination. The primary contribution of my analysis here is to synthesize the literature on the roles of administrative burden and trust in mediating public health policy, and cost factors involved with the utilization of care within the Medicaid population.

Figure 3.2: Self-reported vaccination or intent to vaccinate rates time series by insurance coverage type.



Plotted lines correspond to the percent of individuals surveyed through the Census' Pulse Survey that self-report receiving at least one dose of, or state an intent to receive the COVID-19 vaccine, by provided insurance coverage type. Data is drawn from the raw Pulse Survey for 'Weeks' 22 to 37, corresponding to January through September, 2021; some individuals will report coverage by multiple insurance types, so these measures should not be considered exclusive. Shaded regions correspond to the 95% confidence bounds.

The remainder of the paper is organized as follows. In the next section I offer some context on preventative utilization among the Medicaid population, as well as a summary of the literature examining the associated role of administrative burden. In Section 3 I describe my data and empirical strategy, in Section 4 I report my results, and in Section 5 I offer a discussion of the policy implications and conclude.

3.2 Background

The consistently lower rates of preventative care utilization among Medicaid recipients are well-documented. Research has shown that infants on Medicaid are less likely to receive vaccinations compared to those with private insurance (Hill, Elam-Evans, Yankey, Singleton, & Kang, 2017, 2018). Similarly, the uptake of the HPV vaccine among adolescent Medicaid recipients is lower, as indicated by claims data (Cook et al., 2010), though there is evidence that Medicaid expansions have led to improvements in HPV vaccination rates (Churchill, 2021). When examining influenza vaccinations, a service more akin to COVID-19 vaccinations, rates among Medicaid recipients are significantly lower than

those of the broader population (Stoecker, Stewart, & Lindley, 2017; Naderalvojud et al., 2023). Furthermore, essential preventative services like mammograms and Pap tests show persistently low utilization among low-income women, even after Medicaid expansion, hinting at a systemic issue (Alharbi, Khan, Horner, Brandt, & Chapman, 2019). This is corroborated by findings of later-stage cervical cancer diagnoses in Medicaid recipients (O'Malley, Shema, Clarke, Clarke, & Perkins, 2006).

In light of these patterns, it's not surprising that during the COVID-19 pandemic, vaccination rates among the Medicaid cohort have been markedly lower than the general population. Although a comprehensive data source for overarching trends is lacking, a 2022 study by the Kaiser Family Foundation using data from five states found that Medicaid recipients were 15-20% less likely to have received at least one dose of a COVID vaccine compared to the general populace (Galewitz, 2022). Further, a collaborative report from the National Academy for State Health Policy and Duke-Margolis Center for Health Policy highlighted vaccine hesitancy that mirrors the broader low-income and minority group trends. Concerns over these findings led some lawmakers to attribute the disparities to administrative and technological challenges, prompting a request for the United States Government Accountability Office to investigate these low vaccination rates and recommend enhancements in data sharing and outreach (Casey & Wyden, 2022).

In an attempt to better illustrate these Medicaid vaccine hesitancy trends, I collect data from the CDC for the percent of county populations recorded as having completed the primary series (initial and first booster dose) of COVID-19 vaccinations and plot the time series for these in Figure 3.1. I distinguish between counties within the bottom and top quartiles for percentages of population enrolled in Medicaid in an effort to identify varying propensities between these groups. One can observe that among counties in the top quartile (at or above the 75th percentile for enrollment), vaccination rates are substantially higher. However, it's crucial to note that these raw figures don't adjust for variations in population characteristics or demographics. Given that counties with high Medicaid participation rates often differ significantly from those with low rates, directly comparing these vaccination trends might not provide an accurate picture. As such, I

additionally employ an entropy balancing protocol to compose a weighted average of 25th percentile counties that best approximates the 75th percentile counties on observable demographics (more elaboration on this methodology is provided in Methods section below) and plot this more appropriate comparison as the dashed line. When compared to this balanced comparison group, counties with high Medicaid participation demonstrate substantially lower vaccination rates that persisted over the entire observed time frame.

The reasons behind this hesitancy are likely complex and may largely stem from Medicaid recipients' distinct perceptions of the costs associated with using care services. Research exploring the potential costs tied to enrolling in or obtaining care through Medicaid generally categorizes these costs into two types: implicit and explicit. The implicit costs of the Medicaid enrollment process have been attributed to administrative burden individuals face in receiving admission to the program. Evidence suggests that the time and labor costs of completing paperwork, making phone calls, and similar activities needed to navigate enrollment bureaucracy have substantial impacts on overall program enrollment figures. Fox, Stazyk, and Feng (2020) exploit variations across states in the extent of administrative burden (defined as a composite index based on real-time eligibility, digital access, enrollment ease, and renewal ease) easing for Medicaid enrollment that occurred following the enactment of the ACA. They find a 3% increase in Medicaid enrollments when comparing states which most aggressively pursued administrative burden easing against those which engaged in relatively little. Similarly, Stuber and Kronebusch (2004) find that the perceived poor treatment by officials when attempting to enroll or take advantage of Medicaid, enrollment barriers, and a lack of transparency regarding Medicaid rules all significantly decrease the likelihood of individuals to enroll in the program. From an alternate perspective, Moynihan et al. (2016) examine state-by-state variations in Medicaid enrollment administrative burden finding that policymakers are aware of the barriers this places on potential Medicaid recipients and actively ease or expand these burdens according to political preferences.

For those individuals that do prevail and enroll with their state Medicaid programs, further implicit costs in the form of care-provider discrimination present themselves.

Numerous analyses of health interview surveys (Thorburn & De Marco, 2010; Weech-Maldonado, Hall, Bryant, Jenkins, & Elliott, 2012; Han, Call, Pintor, Alarcon-Espinoza, & Simon, 2015; Alcalá & Cook, 2018; Alcalá, Ng, Gayen, & Ortega, 2020) have repeatedly demonstrated significantly higher likelihoods of perceived discrimination for Medicaid recipients from care-providers when compared to those who are privately-insured. It is unclear whether this perceived discrimination translates to measurable changes in care utilization (Thorburn and De Marco (2010) find no significant decrease in the receipts of prenatal care among a sample of California Medicaid recipients, despite substantially higher self-reporting of discrimination when receiving said care), but the psychic costs of discriminatory treatment are almost certain to influence patients' perceptions of care quality.

The literature on the explicit costs of enrolling or receiving Medicaid has focused primarily on changes in enrollment following adjustments to cost-sharing policies or extent of coverage. Survey data suggests that states which offer more extensive benefits to their Medicaid recipients boast substantially higher enrollment rates (Stuber & Kronebusch, 2004). Other surveys of Medicaid recipients have found that financial concerns (e.g. "worry pay more than expect," or "worry insurance won't cover care") were cited by 65% of the survey participants as perceived barriers to receiving care (E. M. Allen et al., 2017). Wright et al. (2005) examine changes to the enrollments and care utilization among Oregon Medicaid recipients following an increase in cost-sharing in 2003. Of the 1,378 participants in their survey, 44% ultimately exited the program after the cost-sharing increase went into effect; and among those individuals that departed from the Medicaid program and reported "zero income," 68.2% cited the cost increases as the principal reason for unenrolling. As well, they find that for those that unenrolled due to cost concerns, the majority reported having some degree of unmet care in the following year.

It is noteworthy that of these preceding Medicaid cost studies, few if any have explicitly observed *decreases* in care utilization based on Medicaid status alone. Indeed, a voluminous literature has emerged in the past decade examining the effects of Medicaid expansions under the Affordable Care Act on utilization (see Mazurenko et al. (2018) or Antonisse

et al. (2019) for exhaustive literature reviews) and has consistently found recipients consume more care following enrollment. These outcomes are intuitive: as barriers to care diminish through the extension of publicly furnished health insurance, utilization should increase to offset previously unmet care demand. Moreover, as access to care increases, so too should institutional familiarity among the newly-insured. This underscores a beneficial aspect where Medicaid enrollment is anticipated to enhance vaccination tendencies. When combined with the previously discussed administrative burden of enrollment, we would anticipate Medicaid enrollees to show improved vaccination rates compared to the uninsured, yet experience significant reductions when contrasted with the privately insured. Consequently, the overall impact within the general population remains uncertain and a matter for empirical inquiry.

3.3 Empirical strategy

3.3.1 Data

The primary data employed in the analysis here is drawn from the U.S. Census Bureau's Household Pulse Survey, which contains the pool of 17 'weeks'³⁰ of surveys conducted between early January through September, 2021. While subject to certain limitations- in particular relatively low participation rates- these data have been employed successfully in similar settings to examine how the vaccination decision has translated to delayed receipt of other care services (Aslim, Fu, Liu, & Tekin, 2022) and reduced demand for Supplemental Nutrition Assistance Program (SNAP) benefits (Aslim, Fu, Tekin, & You, 2023). These microdata include a rich variety of individual and household demographic variables. Of critical importance is the existence of a public health insurance status variable that permits us to identify Medicaid enrollees that exist within the sample. Since the 'Week 22' period that collected data between January 6 through January 18, 2021, these surveys have solicited participants for their current COVID-19 vaccination status

³⁰These 'week' survey periods cover twelve days (Wednesday to Monday) over which surveys are conducted. Over most of the observed period, the next 'week' would commence the Wednesday two days after the end of the preceding 'week'.

and intentions towards receiving a vaccine; most significantly, for those individuals that indicate no intent or a low likelihood of vaccination, they are further inquired of their reasoning for hesitancy. Prior to exhaustive cleaning, I plot the raw vaccination rate time series in Figure 3.2, grouped by insurance coverage type. After cleaning, I retain 93,823 unique observations with which to perform this analysis. These cleaned data represent all individuals who self-report having coverage through their employer (Group), some other private insurer (Individual), Medicaid, or as having no health insurance. Additionally, all individuals over the age of 65³¹, with household income over \$63,000 or more than 200% of their state's maximum income cap for Medicaid eligibility (I elaborate more on this measure below), and with self-reported coverage through Medicaid and some other source (e.g. those who report having coverage through both Medicaid and Medicare) are also dropped. The summary statistics for a selection of demographic variables can be found in Table 3.1 and are grouped by insurance coverage type.

As is evident in Table 3.1, the Medicaid and uninsured populations are disproportionately unemployed and in receipt of SNAP benefits compared to the privately insured, while also substantially less likely to report either being vaccinated or having an expressed intent to receive the COVID-19 vaccine. This is made all the more stark in Figure 3.2, which illustrates persistently lower vaccination rates (in-line with the vaccination rate trends from Figure 3.1) among the Medicaid population over the entire observed time frame. Taken together, these summary statistics provide additional evidence of vaccine hesitancy within the Medicaid population, as well as highlighting the need for concretely identifying the appropriate counterfactual to Medicaid enrollment.

This analysis delves into the policy implications of Medicaid enrollment's impact on vaccination decisions by exploring two specific dimensions: the effect of Medicaid benefits compared to being uninsured, and the influence of these benefits relative to private insurance coverage (either group or individual). Despite historically lower rates of utilization for other preventive care services among Medicaid enrollees versus the privately insured, Medicaid expansions have shown to enhance utilization for those previously

³¹All survey data is collected from adults ages 18 and older.

uninsured. This investigation addresses two critical angles: whether Medicaid policies should focus on boosting utilization or expanding coverage. Prior research has typically concentrated on one of these aspects, but this study aims to assess both, along with the overall effects, by analyzing three distinct sample groups: one comprising individuals who are either uninsured or solely Medicaid-insured, another limited to those with Medicaid or other insurance types (excluding Medicaid-only), and a final group encompassing all individuals from the refined sample.

In addition to these Census data, I also employ a number of others to supplement my basic analysis. To account for potential endogeneity issues, I utilize state-level Medicaid income eligibility guidelines for the year 2021 collected from the Centers for Medicare and Medicaid Services to estimate each observed household's proximity to the income cap. This percent of Medicaid income eligibility cap is estimated as the ratio of self-reported household income, and the state-of-residence's maximum Medicaid eligibility income cap based on the household dependents-conditional federal poverty level. State Medicaid agencies has discretion to set these income caps with respect to not only federal poverty level, but also conditional on the composition of the household (for instance, requiring that recipients be parents to children that also within the household). I estimate this by first deriving the maximum permitted income for each household based on the state or residence, number of children, and household size, then divide this by either the mid-point or maximum of the self-reported income-bin from the Pulse Survey. Ideally, this measure can capture exogenous cross-state variation in Medicaid eligibility policies that should not directly affect vaccination outcomes, except through changes to Medicaid coverage status. To provide measurements for variations in administrative burden, cost-sharing, and benefits, I employ the Health Affairs Medicaid Accessibility Index (Fox, Feng, Zeitlin, & Howell, 2020), which includes standardized performance scores from 0 (poor) to 1 (excellent) for every state across each of these aspects. The administrative burden index, which draws on data from the Kaiser Family Foundation's 50-state Survey of Medicaid and CHIP Eligibility, Enrollment, and Cost Sharing Policies Reports from 2000 to 2018, examines the differences in how states manage enrollment and renewal barriers for Medicaid and

CHIP, alongside their initiatives to simplify these processes. Notable policy improvements facilitating easier access and maintenance of coverage include the removal of asset tests, transition from in-person to phone or online applications, shortened waiting times for enrollment and renewal, and the option for applicants to self-report their income. The benefit and co-pay coverage index is built on the KFF’s Medicaid Benefits Database which is derived from surveys conducted with Medicaid directors from all 50 states and D.C., focusing on the benefits available to adult beneficiaries, the associated cost-sharing requirements, and benefit limitations. States were scored based on the presence of specific benefits and the absence of co-pays. I additionally collect per capita Medicaid expenditures by state and age-group for 2019 (so as to avoid any unusual policy changes initiated during the COVID-19 pandemic) from the Kaiser Family Foundation website³².

Table 3.1: Summary statistics for cleaned Pulse Survey data, by insurance coverage type.

Statistic	<i>Insurance coverage sample:</i>				
	Group (<i>n</i> = 43, 253)	Individual (<i>n</i> = 15, 843)	Medicaid (<i>n</i> = 28, 241)	Uninsured (<i>n</i> = 14, 176)	Pooled Sample (<i>n</i> = 93, 823)
HH Size	3.223	3.002	3.163	3.169	3.151
Age	42.410	45.363	43.763	42.670	43.205
Is Male	0.249	0.274	0.181	0.323	0.241
Is Married	0.446	0.440	0.266	0.333	0.370
Is Hispanic	0.144	0.139	0.154	0.254	0.159
HH Log Income	10.481	10.344	9.910	10.101	10.233
Is or Intends to Vaccinate	0.741	0.749	0.622	0.599	0.688
HH # of Children	0.884	0.721	1.065	0.879	0.906
SNAP Recipient	0.076	0.075	0.512	0.175	0.221
Is Employed	0.792	0.651	0.402	0.517	0.616
Had COVID	0.152	0.140	0.133	0.140	0.141

All data described above is drawn from the U.S. Census Bureau’s Household Pulse Survey, which contains the pool of 17 weeks of surveys conducted between early January through September, 2021. These cleaned data represent all individuals who self-report having coverage through their employer (Group), some other private insurer (Individual), Medicaid, or as having no health insurance. Additionally, any individuals over the age of 65, with household income over \$63,000 or more than 200% of their state’s maximum income cap for Medicaid eligibility, and with self-reported coverage through Medicaid and some other source (e.g. those who report having coverage through both Medicaid and Medicare) are dropped.

³²Accessible here: <https://www.kff.org/medicaid/state-indicator/medicaid-spending-per-full-benefit-enrollee/>

3.3.2 Methods

To estimate the marginal influence of Medicaid enrollment status on likelihood of vaccinating, I estimate the following reduced-form linear probability model:

$$y_{i,t} = \alpha + \beta \text{MEDICAID}_i + X_i\delta + \phi_{s,t} + \epsilon_{i,t} \quad (3.1)$$

where y_i is an indicator for individual i 's self-reported receipt of a COVID-19 vaccine, or stated intention to receive a full vaccine treatment regimen; MEDICAID_i is the treatment, an indicator variable for whether the observed individual is insured through Medicaid; X_i is a vector of individual- and household-level demographic controls; and $\phi_{s,t}$ are state by week fixed effects. Incorporating state-by-week fixed effects allows for a non-parametric, flexible approach to adjust for broader trends in vaccine uptake within each state. Specifically, these fixed effects should control for any spatiotemporal fluctuations in state-specific policies that could influence the availability or accessibility of COVID-19 vaccines. The coefficient of interest then is β , which should capture the marginal effect that Medicaid enrollment status has on the likelihood of vaccinating.

An immediate concern in interpreting any results of (3.1) as causal is that the selection into the 'treatment' group is not random and likely endogenous with other unobserved variables. For instance, consider the notable absence of the political leanings of observed individuals within the demographic data: recent studies have demonstrated the influence of political party-alignment on vaccine hesitancy (Hornsey et al., 2020; Fridman et al., 2021), while earlier surveys have found evidence of liberal political preferences among the Medicaid population based on the high degrees of correlation between Democratic voter registration (Baicker et al., 2019) and vote shares (Hollingsworth et al., 2019) following Medicaid expansions; this would suggest caution is warranted when attempting to perform any causal inference based on (3.1) alone, as unobserved political persuasion is likely confounding any coefficient estimates.

It is worth noting that the most plausible confounding factors are likely to bias the

coefficient estimates positively. If it were the case that individuals who self-identify as conservatives exhibited decreased likelihoods for both vaccinating and enrolling in Medicaid, this would imply the estimated marginal effects from (3.1) are inflated positively. Similarly, it may be the case that Medicaid recipients are more sickly at the margin relative to those who are Medicaid eligible but uninsured, and the lack of inclusion of individual-level health status indicators may bias the estimated treatment effect; but even then, preliminary evidence (Soares et al., 2021) suggests the presence of comorbidities induces a greater proclivity to vaccinate, and as such, a more-infirm Medicaid population should exhibit greater likelihoods overall, thus biasing the coefficient upwards. Taken together with my previous assertion that- absent the discouraging effects of principal interest here from enrolling in or receiving care through Medicaid- Medicaid recipients should presumably demonstrate a higher likelihood of vaccinating based on increased institutional familiarity, at least when comparing the uninsured and Medicaid populations. When comparing the Medicaid and privately insured populations, the potential influence of this bias is less clear *a priori*: Medicaid recipients may still be more sickly (for instance, if they are disabled and unable to receive insurance through an employer), but could also have more limited mobility which prevents easy access to vaccination services.

To derive meaningful estimates for the magnitude of these specific costs of Medicaid enrollment and utilization, I apply an entropy balancing method (Hainmueller, 2012) to address potential endogeneity in the $MEDICAID_i$ ‘treatment’ variable. Entropy balancing is a preprocessing technique that assigns weights to control group observations to ensure that the distributions of covariates between the ‘treatment’ and control groups are equivalent, thus “balancing” the covariates. This process aims to balance all relevant covariates at once, emulating the conditions of a randomized experiment and enabling the derivation of more credible effect estimates. When successful, entropy balancing ensures that differences in these covariates do not introduce bias into the results, enhancing the validity of the causal inferences. This method is an extension of traditional propensity score weighting, which typically uses logistic regression to calculate weights and then conducts balance checks to confirm the equalization of covariate distributions. Critically, entropy

balancing requires fewer assumptions to produce credibly causal estimates, and is less vulnerable to bias when misspecified compared to propensity score methods (Hainmueller, 2012). Indeed, simulation studies have demonstrated entropy balancing as outperforming more *ad hoc* propensity score approaches (Zhao & Percival, 2017). Nonetheless, as a robustness check I also estimate a similar propensity score matching model and provide details on this and the results in appendix 3.7.1.

To provide more detail on the specifics involved with estimating these entropy balancing models: This approach involves recalculating the LPM outlined in equation (3.1), applying a distinct weighting to each observation based on the vector $\mathbf{w} = (w_1, w_2, \dots, w_{n_0})$. Here, w_i is set to 1 for units where $\text{MEDICAID}_i = 1$, and for those where $\text{MEDICAID}_i = 0$, w_i is chosen by the following reweighting scheme:

$$\min_{w_i} H(w) = \sum_{\{i|\text{MEDICAID}=0\}} h(w_i) \quad (3.2)$$

subject to balance constraints:

$$\begin{aligned} \sum_{\{i|\text{MEDICAID}=0\}} w_i X_i &= \sum_{\{i|\text{MEDICAID}=1\}} X_i, \quad \text{or} \\ \sum_{\{i|\text{MEDICAID}=0\}} w_i (X_i - \mu_i)^r &= \sum_{\{i|\text{MEDICAID}=1\}} (X_i - \mu_i)^r \quad \text{with } r \in 1, \dots, R \end{aligned}$$

where $h(\cdot)$ is the entropy distance metric, $h(w_i) = w_i \log(w_i/q_i)$, μ_i the sample mean for either the control or Medicaid groups, and R is the number of balance constraints imposed on the covariate moments of the reweighted control group. For the purposes of my estimates, I set base weights as uniform, $q_i = 1/n_0$, and $R = 3$ so that derived control weights satisfy two appealing conditions: The loss function described by equation (3.2) penalizes larger deviations from a uniform weighting scheme, ensuring no excessive over-weighting of a few control observations; and for the selected covariates that samples are balanced on, the mean, variance and skew for the distributions of said covariates are identical between the Medicaid and control groups. I balance on covariates for gender, race,

ethnicity, marital status, logged household income, household size, number of children, age, current housing situation (renting, owned, cohabitating with family, etc.), and SNAP benefit status. In a similar manner to the naive OLS regressions, all entropy balanced models are estimated across three sub-samples to identify differing responses across the Medicaid vs. uninsurance, and Medicaid vs. private insurance margins, as well as net pooled effects.

In addition to these direct entropy balancing models, I also leverage the computed percent of Medicaid income eligibility cap measure in two different methodological settings in an attempt to derive more credibly causal estimates. In the first of these alternative approaches, I use this percentage measure as a running variable and the 100% level as a cutoff in a sharp regression discontinuity design (RDD)³³. Although eligibility for Medicaid is frequently affected by various unobserved factors, surpassing the maximum income threshold generally renders most households ineligible unless they qualify for a specific exemption, such as one related to a medical disability. Because a primary source of bias in naive OLS estimates could stem from the disproportionate number of individuals with poor health within the Medicaid population, utilizing this income threshold provides a strategic method to detect any marginal changes in vaccination practices among Medicaid recipients of more typical health status. If some data preprocessing method is employed to balance observables between control and Medicaid groups, these RDD estimates should be able to identify causal marginal effects of enrollment status as households and individuals on either side of this cutoff are approximately identical, notwithstanding Medicaid eligibility. Moreover, a household or individual's location immediately to the left or right of the cutoff should be fairly random and subject to cross-state variation in eligibility limits. This can be thought of as a generalization of the method used by H. Allen et al. (2021), which leverages the 138% of the federal poverty level eligibility cutoff on a propensity score balanced sample in an RDD setting.

³³As an example, if you lived in a state where the maximum Medicaid eligible income was 100% of the FPL, then the 100% cutoff would represent the FPL for your household size and number of dependents. Alternatively, if you lived in an expansion state where the maximum eligible income was 138% of the FPL, then the 100% cutoff would represent 138% of the FPL for your household size and number of dependents.

In the second alternative to the primary entropy balancing results, I employ the Medicaid income eligibility cap measure as a synthetic instrumental variable within a two-stage least squares (2SLS) model to account for any biasing selection into ‘treatment’. If the inclusion of direct measures for income and household size as covariates are sufficiently capturing the influence of these factors on the vaccination decision, then the simultaneous inclusion of the income eligibility cap should describe only (and comparatively random) cross-state variation in eligibility guidelines. If this assumption is correct, that would satisfy the necessary exclusion restriction for this income cap measure as it should only affect the vaccination decision through Medicaid enrollment status (and consequently uncorrelated with the second-stage error term). Given that the income cap measure has a clear correlation with the endogenous Medicaid regressor, this indicates its potential as a valid instrument.

Lastly, to produce evidence of the relative impact of the suggested mechanisms, I propose (3.3) as an amended version of (3.1)

$$y_{i,t} = \alpha_3 + \beta_1 \text{MEDICAID}_i + \beta_2 \text{MEDICAID}_i \times \text{PC}_s + \beta_3 \text{MEDICAID}_i \times \text{BC}_s + \beta_4 \text{MEDICAID}_i \times \text{AB}_s + X_i \delta_3 + \phi_{s,t} + \epsilon_{3i,t} \quad (3.3)$$

this includes interaction terms between MEDICAID_i and per capita Medicaid expenditures (PC_s), the benefit and copay coverage index (BC_s), and the administrative burden index (AB_s) compiled by Fox, Feng, et al. (2020). If my administrative burden mechanism hypothesis is correct, I should observe positive coefficients for β_2 , β_3 , and β_4 , suggesting that as per capita spending increases, and copays and administrative burden decrease (an increase in the index score), the likelihood of vaccinating increase. To address potential collinearity concerns between these mechanism variables, I estimate multiple models that variously include the different interaction terms. As well, to specifically identify tangible administrative burden effects, I utilize data collected on the ease of the Medicaid application process from Code for America (2019) to observe any categorical differences in the vaccination decision within the Medicaid population between states with more and less-onerous enrollment procedures.

3.4 Results

I estimate the naive OLS regressions according to the Linear Probability Model specified by equation (3.1). Results are listed in Table 3.2. Estimates in column one describe the marginal change in vaccination likelihood of Medicaid status compared to the uninsured population; column three estimates have the same interpretation, but compare vaccination outcomes between the Medicaid and privately-insured populations; column two estimates these changes for the pooled uninsured, Medicaid, and privately-insured populations. All three estimates include individual-level covariates, as well as state-week and state-MSA (if surveyed individual resides in one of the fifteen largest MSAs in the US) fixed effects. Note that because the monetary cost for the COVID-19 vaccine at this time was zero, the extent of cost-sharing associated with an individual's insurance coverage should not affect the vaccination decision. I observe that on average, Medicaid recipients demonstrate a 4.1 percentage point higher propensity to vaccinate, relative to those who are not insured, and that these marginal effects are highly statistically significant. Conversely, when compared to the privately insured (those with group or individual private insurance), the Medicaid population appears to be exhibiting a lower likelihood to be vaccinated of 6.3 percentage points. In the pooled sample, Medicaid recipients remain 3.4 percentage points less likely to vaccinate, suggesting that in absolute terms when considering the relative prevalence of private insurance coverage and uninsurance, Medicaid recipients appear to be less likely to receive the COVID-19 vaccine than the general population of similar households. This would appear to justify my line of inquiry here, suggesting that- at least relative to the privately insured- some other underlying aspect of the Medicaid experience is categorically discouraging recipients from electing to receive vaccination.

One concern already noted with these naive OLS results is that there may be some selection into treatment effects that are biasing estimates. To address this I employ an entropy balancing approach according to equation (3.2) to derive weights for control observations that balance covariate distributions with the Medicaid recipients, then re-estimate the LPM specified under equation (3.1) with the inclusion of these observational weights. I

Table 3.2: Naive OLS estimates for likelihood of receiving or intention to receive vaccine by insurance coverage type.

<i>Sub-sample:</i>	<i>Dependent variable:</i>		
	Vaccine Status		
	Medicaid Only & Uninsured	Pooled	Medicaid Only & Insured
Medicaid	0.041*** (0.005)	-0.034*** (0.004)	-0.063*** (0.004)
<i>Fixed-effects</i>			
Week×State	Yes	Yes	Yes
Week×MSA	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	42,417	93,823	79,647
R ²	0.18096	0.16820	0.17634
Within R ²	0.07445	0.07972	0.08851

Clustered (Week×State) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

All coefficient estimates are derived from a fixed effects linear probability model and represent the marginal change in likelihood of being vaccinated, or having a stated intent to vaccinate by insurance type. Estimates in column one describe the marginal change in vaccination likelihood of Medicaid status compared to the uninsured population; column three estimates have the same interpretation, but compare vaccination outcomes between the Medicaid and privately-insured populations; column two estimates these changes for the pooled uninsured, Medicaid, and privately-insured populations. Covariates include: individual age, sex, marriage status, ethnicity, race, logged household income, home ownership status, household size, number of children in household, employment status, whether they had contracted COVID at some point, and whether they had difficulty with paying bills. Estimates include state-week and MSA-week fixed effects; all standard errors are clustered at the state-week level.

perform these weighting and estimation procedures for each of the three sub-samples, and list the results in panel (a) of Table 3.3. As a robustness check, in panel (b) I use the same entropy balancing weights for each respective sub-sample, but employ them using a weighted fixed effects logistic regression estimator instead of the LPM approach.

Comparing the entropy balanced results in Table 3.3 against those in Table 3.2, I observe virtually identical results. Relative to the uninsured, the Medicaid population is exhibiting a 3.4 to 3.6 percentage point higher likelihood to vaccinate, depending on whether using the LPM or logit results. This would appear to corroborate the claim that more consistent or easier access to care afforded by Medicaid coverage, as well as the development of greater institutional familiarity from being enrolled in the program has translated to a greater proclivity to vaccinate against COVID. When weighted to ensure the most demographically-similar control group within both the pooled and insured (private or Medicaid) populations, I observe a significantly negative estimates and interpret this as Medicaid enrollees exhibiting 3.1 and 6.9 percentage points lower likelihoods of vaccinating (using the LPM estimates). These results appear to align with my primary argument that the barriers to enrollment and care utilization, and the general sense of mistrust unique to the Medicaid population are discouraging vaccine uptake at rates as high as other conventional insurance populations (this is tested directly in a later analysis).

The interpretations of the results in rows one and three of Table 3.3 are relatively straightforward, showing the marginal effects of Medicaid enrollment on vaccination rates compared to the uninsured and privately insured groups, respectively. However, interpreting the pooled results necessitates further clarification. Similar to the naive OLS findings, the pooled results offer a weighted average of the effects within the two sub-populations, reflecting the general population's distribution of private insurance and uninsurance. In the context of entropy balancing, this concept is extended to consider that an ideal control group for comparison should include both uninsured and privately insured individuals to address selection biases. This method also leverages the advantage of a larger sample size, reducing the impact of atypical members within the uninsured (e.g., young, healthy individuals who forego insurance) and privately insured groups, thereby

enhancing the representativeness of the results.

Table 3.3: Entropy balanced estimates for likelihood of receiving or intention to receive vaccine by insurance coverage type.

<i>Sub-sample:</i>		<i>Dependent variable:</i>		
		Vaccine Status		
		Medicaid Only & Uninsured	Pooled	Medicaid Only & Insured
<i>(a) Linear probability estimates</i>	Medicaid	0.034*** (0.007)	-0.031*** (0.005)	-0.069*** (0.007)
	Observations	42,417	93,823	79,647
	R ²	0.19220	0.19744	0.22487
	Within R ²	0.08056	0.08544	0.08726
<i>(b) Fixed effects logistic regression estimates</i>	Logit coefficients	0.186*** (0.038)	-0.162*** (0.028)	-0.387*** (0.041)
	Estimated Marginal effects	0.036*** (0.007)	-0.028*** (0.006)	-0.064*** (0.012)
	Observations	42,297	93,734	79,557
	Squared Correlation	0.17132	0.15090	0.15185
	Pseudo R ²	0.17357	0.23276	0.25776
	BIC	57,626.1	104,991.4	86,167.7

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

((a)) All coefficient estimates are derived from a fixed effects linear probability model weighted according to the entropy balancing protocol outlined in equation (3.2). Estimates in column one describe the marginal change in vaccination likelihood of Medicaid status compared to the uninsured population; column three estimates have the same interpretation, but compare vaccination outcomes between the Medicaid and privately-insured populations; column two estimates these changes for the pooled uninsured, Medicaid, and privately-insured populations. Covariates include: individual age, sex, marriage status, ethnicity, race, logged household income, home ownership status, household size, number of children in household, employment status, whether they had contracted COVID at some point, and whether they had difficulty with paying bills. Estimates include state-week and MSA-week fixed effects; all standard errors are clustered at the state-week level.

((b)) All coefficient estimates are derived from a fixed effects logistic regression model weighted according to the entropy balancing protocol outlined in equation (3.2). The estimated marginal effects in the second row have the same interpretation as the primary results outlined in panel (a).

In Tables 3.4 and 3.5 I perform some stratification analysis and re-estimate the three primary entropy balanced LPMs from Table 3.2 across several key subgroups. In Table 3.4 I focus on heterogeneous responses to Medicaid status across individual and household demographics and characteristics. It is noteworthy that across all of these subgroups, the

differences between coverage type sub-samples remain qualitatively the same: Relative to the privately insured, the Medicaid population demonstrates lower propensities to vaccinate, and that these negative responses persist in the pooled populations but are attenuated towards zero. In most subgroups, the Medicaid continues to demonstrate statistically significant higher likelihoods to vaccinate relative to the uninsured population. Intriguingly, SNAP recipients and black individuals enrolled in Medicaid demonstrate null effects in this regard, while Hispanic Medicaid recipients are less likely to vaccinate relative to the uninsured. These findings underscore the significant impact of mistrust in government on vaccination decisions among minority groups, aligning with previous research (Khan, Ali, Adelaine, & Karan, 2021). Additionally, the data related to SNAP participation suggest potential spillovers in health service use linked to SNAP involvement. Notably, the SNAP estimates within both the pooled and insured subsets closely mirror the primary findings from Table 3.3, with a marked contrast observed only when compared with the uninsured group. This discrepancy can likely be attributed to Medicaid enrollees' higher rates of SNAP participation relative to the uninsured. It may also reflect a greater propensity for vaccination among SNAP beneficiaries, possibly due to increased trust in institutions.

In Table 3.5, I categorize the entropy balancing outcomes based on state-specific factors, including Medicaid expansion status and political leanings, indicated by the vote shares from the 2016 presidential election. Since a key aspect of Medicaid expansions involves simplifying enrollment processes (Fox, Feng, et al., 2020), analyzing variations across these policy dimensions can shed light on potential mechanisms at play. Moreover, the effectiveness of state political organizations in facilitating vaccine distribution has differed significantly across ideological lines, making vote shares a useful metric to explore diverse responses in this context. Most noteworthy is that Medicaid recipients in non-expansion states appear to be no more likely to vaccinate, compared to the uninsured, which again further bolsters the administrative burden argument if these states have also not updated or simplified enrollment procedures. Also intriguing is that Trump states appear to have more exaggerated Medicaid enrollment effect estimates across the board, with larger

Table 3.4: Entropy balanced estimates for likelihood of receiving or intention to receive vaccine by insurance coverage type, stratified by individual characteristics.

<i>Sub-sample:</i>	<i>Dependent variable:</i>		
	Vaccine Status		
	Medicaid Only & Uninsured	Pooled	Medicaid Only & Insured
Black individuals	0.0111 (0.0207)	-0.0204 (0.0141)	-0.0498*** (0.0183)
Observations	5,073	9,339	7,754
Hispanic individuals	-0.0267* (0.0140)	-0.0356*** (0.0111)	-0.0360** (0.0144)
Observations	7,942	14,902	11,300
Men	0.0722*** (0.0133)	0.0059 (0.0105)	-0.0531*** (0.0121)
Observations	9,675	22,618	18,045
Women	0.0165** (0.0082)	-0.0430*** (0.0060)	-0.0720*** (0.0081)
Observations	32,742	71,205	61,602
Employed	0.0307*** (0.0101)	-0.0498*** (0.0067)	-0.0725*** (0.0073)
Observations	18,703	57,837	50,501
Unemployed	0.0209** (0.0101)	-0.0201*** (0.0075)	-0.0670*** (0.0108)
Observations	23,714	35,986	29,146
SNAP recipients	0.005 (0.014)	-0.031*** (0.009)	-0.068*** (0.013)
Observations	16,939	20,749	18,262

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

All coefficient estimates are derived from a fixed effects linear probability model weighted according to the entropy balancing protocol outlined in equation (3.2), stratified by self-reported personal or household characteristics. Estimates in column one describe the marginal change in vaccination likelihood of Medicaid status compared to the uninsured population; column three estimates have the same interpretation, but compare vaccination outcomes between the Medicaid and privately-insured populations; column two estimates these changes for the pooled uninsured, Medicaid, and privately-insured populations. Covariates include: individual age, sex, marriage status, ethnicity, race, logged household income, home ownership status, household size, number of children in household, employment status, whether they had contracted COVID at some point, and whether they had difficulty with paying bills. Estimates include state-week and MSA-week fixed effects; all standard errors are clustered at the state-week level.

percentage point increases (decreases) to vaccination likelihood compared to the uninsured (privately insured).

Table 3.5: State stratification entropy balancing analysis by political and expansion policy status.

<i>Sub-sample:</i>	<i>Dependent variable:</i>		
	Vaccine Status		
	Medicaid Only & Uninsured	Pooled	Medicaid Only & Insured
Medicaid Expansion States	0.035*** (0.008)	-0.031*** (0.006)	-0.069*** (0.007)
Observations	38,567	85,017	72,591
Not Medicaid Expansion States	0.022 (0.019)	-0.024* (0.014)	-0.072*** (0.024)
Observations	3,850	8,806	7,056
Trump States in 2016	0.046*** (0.011)	-0.035*** (0.008)	-0.090*** (0.012)
Observations	18,225	41,500	34,943
Not Trump States in 2016	0.029*** (0.009)	-0.026*** (0.007)	-0.056*** (0.008)
Observations	24,192	52,323	44,704

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

All coefficient estimates are derived from a fixed effects linear probability model weighted according to the entropy balancing protocol outlined in equation (3.2), stratified by state-of-residence Medicaid policies and political preferences. Estimates in column one describe the marginal change in vaccination likelihood of Medicaid status compared to the uninsured population; column three estimates have the same interpretation, but compare vaccination outcomes between the Medicaid and privately-insured populations; column two estimates these changes for the pooled uninsured, Medicaid, and privately-insured populations. Covariates include: individual age, sex, marriage status, ethnicity, race, logged household income, home ownership status, household size, number of children in household, employment status, whether they had contracted COVID at some point, and whether they had difficulty with paying bills. Estimates include state-week and MSA-week fixed effects; all standard errors are clustered at the state-week level.

To examine more concretely the motivations for vaccine hesitancy, I restrict the sample to only those individuals who had not been vaccinated and had no stated intent to vaccinate and replicate the entropy balanced specifications from Table 3.2 on a schedule of possible hesitancy reasons³⁴. The results of these regressions are listed in Table 3.6;

³⁴The Pulse Survey only queried participants that provided a vaccination intention response of "probably get a vaccine," "be unsure about getting a vaccine," "probably NOT get a vaccine," or "definitely NOT get a vaccine" for their hesitancy reason. Although there's a potential discrepancy between those who express a willingness to get vaccinated but don't intend to follow through, and those who initially state

these coefficients should be interpreted as the marginal likelihood of reporting a given reason for not vaccinating when enrolled in Medicaid. The results for all hesitancy reasons by each insurance coverage sub-sample with the Medicaid and uninsured population in column one, the Medicaid and privately-insured populations in column three, and the pooled uninsured, Medicaid, and privately-insured populations in column three. The results provide nuanced evidence for the validity of the trust mechanism hypothesis: Medicaid recipients were 2-2.5 percentage points more likely to cite vaccine or government trust concerns as a reason for vaccine hesitancy, but these effects are only observable in the pooled and insured sub-samples. Medicaid recipients were also more likely to report side effects concerns across all sub-samples. Also intriguing is that across all of these sub-samples, Medicaid recipients were less likely to report cost concerns associated with receiving the COVID-19 vaccine. Again, these results do not imply that Medicaid recipients are not offering these reasons for their hesitancy, but that at margin, when compared to similar insured individuals, Medicaid recipients are no more likely to offer these as explanations.

3.4.1 Income eligibility cap analysis

To amend the preceding analyses and provide more plausibly causal estimates for the influence of Medicaid enrollment status on the vaccination decision, I employ a computed income eligibility cap measure within a few different methodological settings. Recall that the percent of Medicaid income eligibility cap is estimated as the ratio of the midpoint of the self-reported household income bin, and the state-of-residence's maximum Medicaid eligibility income cap based on the household dependents-conditional FPL. This metric is designed to capture exogenous variations in Medicaid eligibility policies across states, which should not directly influence vaccination outcomes, except indirectly through changes in Medicaid coverage status.

they're not interested but end up getting vaccinated, my analysis is concentrated on individuals who explicitly expressed no intention to get vaccinated. Despite the self-reported nature of these responses, which may introduce some inaccuracies due to dishonesty, the reasons given by individuals who assertively declare their unwillingness to vaccinate are likely more reliable and indicative of genuine sentiments.

Table 3.6: Entropy balanced likelihood estimates for vaccine hesitancy reason by insurance coverage type.

<i>Sub-sample:</i>	Medicaid Only & Uninsured	Pooled	Medicaid Only & Insured
<i>Dependent Variable</i>			
Side Effects	0.072*** (0.012)	0.048*** (0.009)	0.024* (0.013)
Won't Work	0.011 (0.010)	0.012 (0.008)	0.013 (0.011)
Don't Need	-0.017** (0.008)	-0.002 (0.007)	0.013 (0.009)
Not Recommended	0.017*** (0.006)	0.003 (0.005)	-0.015** (0.008)
Wait and See	0.044*** (0.013)	0.021** (0.010)	0.003 (0.013)
Cost Concern	-0.062*** (0.008)	-0.035*** (0.005)	-0.014** (0.007)
Don't Trust Vaccine	0.015 (0.011)	0.022** (0.009)	0.025** (0.012)
Don't Trust GOV	0.004 (0.010)	0.020*** (0.007)	0.023** (0.010)
COVID Not Threat	-0.005 (0.005)	0.0009 (0.004)	0.003 (0.005)
Observations	16,379	29,300	23,609

Clustered (Week×State) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

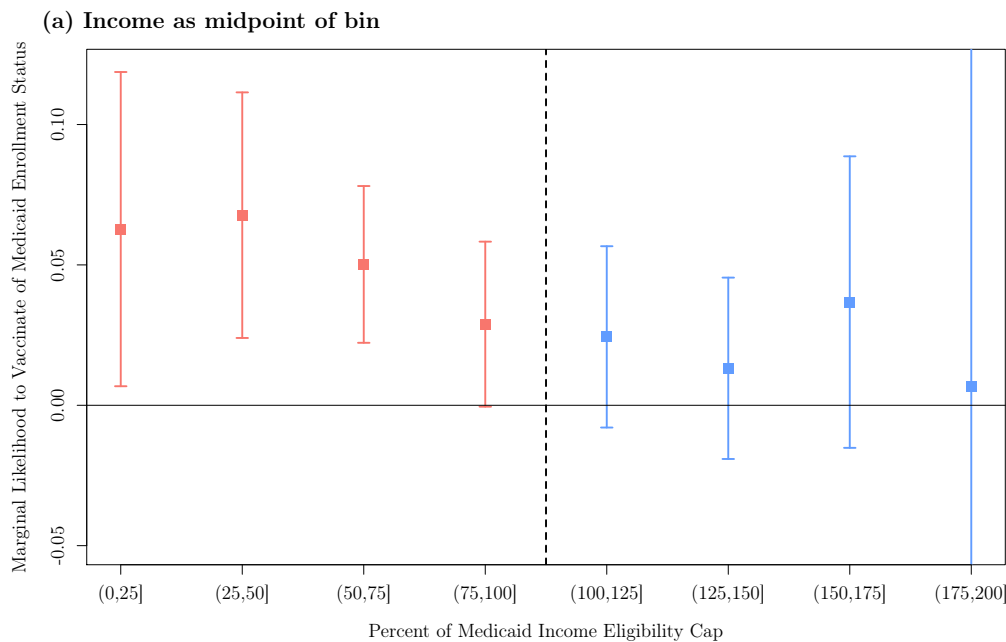
All coefficient estimates are derived from a fixed effects linear probability model weighted according to the entropy balancing protocol outlined in equation (3.2). Results in column one are estimated on the Medicaid and uninsured population sub-sample; column three on the Medicaid and privately-insured populations; and column two is for the pooled uninsured, Medicaid, and privately-insured populations. The provided reasons are, in order: I am concerned about possible side effects of a COVID-19 vaccine; I don't know if a COVID-19 vaccine will protect me; I don't believe I need a COVID-19 vaccine; my doctor has not recommended it; I plan to wait and see if it is safe and may get it later; I am concerned about the cost of a COVID-19 vaccine; I don't trust COVID-19 vaccines; I don't trust the government; I don't think COVID-19 is that big of a threat; and other. All coefficient estimates represent the marginal change in likelihood of citing these reasons for vaccine hesitancy associated with Medicaid enrollment. Covariates include: individual age, sex, marriage status, ethnicity, race, logged household income, home ownership status, household size, number of children in household, employment status, whether they had contracted COVID at some point, and whether they had difficulty with paying bills. Estimates include state-week and MSA-week fixed effects; all standard errors are clustered at the state-week level.

To start, I estimate the same entropy balancing model on the Medicaid and uninsured sub-sample described in the column one of Table 3.3, but stratify by binned percentages of this maximum income cap. A potential issue with considering the Medicaid population as uniform and comparable to either uninsured or privately insured groups is the inclusion of individuals who qualify for Medicaid due to severe illness or disability, despite surpassing income thresholds. Analyzing vaccination behavior changes in relation to the income cap percentage offers a strategy to differentiate between these individuals and those who are typically healthy yet lose eligibility upon exceeding the 100% income threshold. I plot the coefficient estimates for these stratified models in Figure 3.3 against the percent of Medicaid income eligibility cap. Because the Pulse Survey provides binned household income measures, panel (a) uses the midpoint, and panel (b) the maximum value of these bins as the household income from which the percent of cap measure is derived³⁵. One can observe in both panels a clear decline in vaccination propensities when surpassing the 100% income cap threshold, suggesting that conventional, income-eligible Medicaid recipients are more likely to utilize these services with coverage.


In Table 3.7 I estimate a series of models employing local polynomial regression discontinuity designs (RDD), utilizing the percentage of the Medicaid income eligibility cap as the running variable with 100% of the cap serving as the threshold. Following the methodology of H. Allen et al. (2021), these estimates are based on samples matched via propensity scores as described by equation (3.5) and detailed in Appendix 3.7.1. The regression discontinuity models are computed using standard, bias-corrected, and robust estimators as per Calonico et al. (2014, 2015, 2020), applying both linear and quadratic local polynomial regressions with a triangular kernel and an MSE-optimization bandwidth selector. Surprisingly, despite the clear trends observed from Figure 3.3, none of these RDD models yield significant point estimates for this eligibility cutoff. To make sense of these results, I illustrate one RDD by plotting the local polynomial (LOESS) regressions on the propensity score matched Medicaid recipients and uninsured sub-sample in panel (a) of Figure 3.4. Although there might be a noticeable increase at the cutoff point,

³⁵When using the max of bin as income, there are no households at or below 25% of the Medicaid income cap, so this coefficient is omitted due to lack of data in panel (b).

Figure 3.3: Entropy balanced results for Medicaid recipients against uninsured sub-sample, stratified by percent of Medicaid income eligibility cap.



((a)) Entropy balancing coefficient estimates represent Medicaid-status' effect on likelihood to vaccinate, stratified by distance from estimated income eligibility cap. Percent of Medicaid income eligibility cap is estimated as the ratio of the midpoint of the self-reported household income bin, and the state-of-residence's maximum Medicaid eligibility income cap based on the household dependents-conditional FPL. Models are estimated on the Medicaid and uninsured population sub-sample. Error bars represent the 95% confidence bounds.



((b)) Entropy balancing coefficient estimates represent Medicaid-status' effect on likelihood to vaccinate, stratified by distance from estimated income eligibility cap. Percent of Medicaid income eligibility cap is estimated as the ratio of the maximum of the self-reported household income bin, and the state-of-residence's maximum Medicaid eligibility income cap based on the household dependents-conditional federal poverty level. Models are estimated on the Medicaid and uninsured population sub-sample. Error bars represent the 95% confidence bounds.

significant heteroskedasticity is present, as shown by the wide confidence intervals in the left-hand regression. An uneven distribution of observations around the cutoff could challenge the validity of the RDD approach. Specifically, in panel (b), when plotting a histogram for the running variable, there's a notable clustering of observations just beyond the 100% threshold. Despite the theoretical merits of this method, the Pulse Survey's income bin data may lack the precision required for effective RDD application, especially when contrasted with the more detailed household eligibility data from state Medicaid agencies highlighted by H. Allen et al. (2021).

As an alternative to the RDD approach, I additionally estimate the naive LPM model described by equation (3.1), but using the percentage of Medicaid income eligibility cap in the first stage as an instrument to account for endogenous selection into Medicaid participation. The results for this 2SLS are listed in Table 3.8. Across all three sample specifications, the income eligibility cap appears to be a strong instrument and produces sensible first-stage results: As expected, the likelihood of being enrolled in Medicaid declines as the percentage of the income cap increases. The second-stage results are somewhat similar to the naive OLS and entropy balanced results, but differ in a few key ways. First, after accounting for selection effects with this income cap instrument, Medicaid recipients appear to be no more likely to vaccinate compared to the uninsured. Secondly, within both the pooled and insured samples, the marginal effects of Medicaid enrollment status are substantially larger than in the primary estimates (though these estimates are only significant at the 10% level for the insured sub-sample). Collectively, these findings offer a less optimistic view of the impact of Medicaid enrollment on vaccine uptake. When adjusting for selection biases, the advantages of institutional support and acquired knowledge seem to be minimal, whereas the deficiencies in care quality and accessibility, especially in comparison to what is available to the privately insured, become significantly more evident.


Table 3.7: Regression discontinuity results on propensity score matched samples, percent of Medicaid income eligibility cap as running variable.

<i>Sub-sample:</i>	<i>Dependent variable:</i>					
	Vaccine Status					
	Medicaid Only & Uninsured		Pooled		Medicaid Only & Insured	
	<u>Linear</u>	<u>Quadratic</u>	<u>Linear</u>	<u>Quadratic</u>	<u>Linear</u>	<u>Quadratic</u>
Conventional	-0.006 (0.043)	-0.017 (0.042)	0.034 (0.041)	-0.015 (0.031)	0.008 (0.038)	-0.020 (0.032)
Bias-Corrected	-0.015 (0.043)	-0.033 (0.042)	0.026 (0.041)	-0.026 (0.031)	0.002 (0.038)	-0.026 (0.032)
Robust	-0.015 (0.045)	-0.033 (0.045)	0.026 (0.043)	-0.026 (0.032)	0.002 (0.040)	-0.026 (0.035)
Observations	28,352	28,352	56,482	56,482	56,482	56,482
BW est.	13.534	36.250	15.043	27.660	8.536	27.474
BW bias	29.766	48.731	32.615	48.049	21.995	38.678
ρ	0.455	0.744	0.461	0.576	0.388	0.710

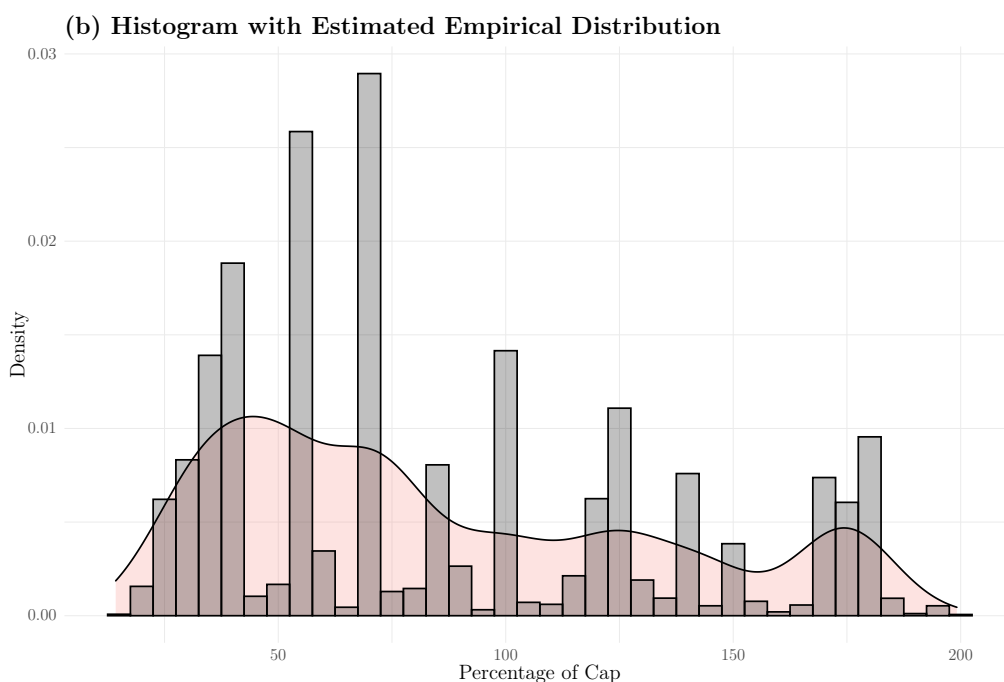
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Coefficient estimates represent local polynomial regression discontinuity point estimates using the percentage of Medicaid income eligibility cap as the running variable and 100% of cap as the cutoff. Outcome variable, vaccination status, is an indicator for self-reported receipt or intent to receive the COVID-19 vaccine. Following H. Allen et al. (2021), these estimates are performed on the propensity score matched samples derived according to equation (3.4) and the procedure outlined in Appendix 3.7.1. Percent of Medicaid income eligibility cap is estimated as the ratio of the midpoint of the self-reported household income bin, and the state-of-residence's maximum Medicaid eligibility income cap based on the household dependents-conditional federal poverty level. Discontinuity models are estimated using conventional, bias-corrected and robust estimators according to Calonico et al. (2014, 2015, 2020), with both linear and quadratic local polynomial regressions employing a triangular kernel and a MSE-optimization bandwidth selector.

Figure 3.4: Regression discontinuity analysis on propensity score matched samples (Medicaid recipients against uninsured).



((a)) Plotted local polynomial (LOESS) regressions on the propensity score matched Medicaid recipients and uninsured sub-sample derived from equation (3.4), with percentage of Medicaid income eligibility cap as the running variable. Outcome variable, vaccination status, is an indicator for self-reported receipt or intent to receive the COVID-19 vaccine. Percent of Medicaid income eligibility cap is estimated as the ratio of the midpoint of the self-reported household income bin, and the state-of-residence's maximum Medicaid eligibility income cap based on the household dependents-conditional federal poverty level. Shaded regions represent the 95% confidence bounds.



((b)) Histogram and empirical distribution of observations within the propensity score matched Medicaid recipients and uninsured sub-sample, on household percentage of the Medicaid income eligibility cap.

Table 3.8: 2SLS estimates using percent of income cap as instrument.

<i>Sub-sample:</i>	<i>Dependent variable:</i>		
	Vaccine Status		
	Medicaid Only & Uninsured	Pooled	Medicaid Only & Insured
<i>First Stage</i>			
% of Income Cap	-0.0026*** (0.0002)	-0.0018*** (0.0001)	-0.0018*** (0.0001)
F-test (IV only)	240.82	310.66	301.69
Wald (IV only), p-value	3.43×10^{-47}	7.01×10^{-64}	1.44×10^{-59}
<i>Second Stage</i>			
Medicaid	0.0003 (0.0668)	-0.1504** (0.0614)	-0.1188* (0.0612)
<i>Fit statistics</i>			
Observations	42,417	93,823	79,647
R ²	0.17971	0.15928	0.17435
Within R ²	0.07305	0.06986	0.08630

Clustered (Week×State) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

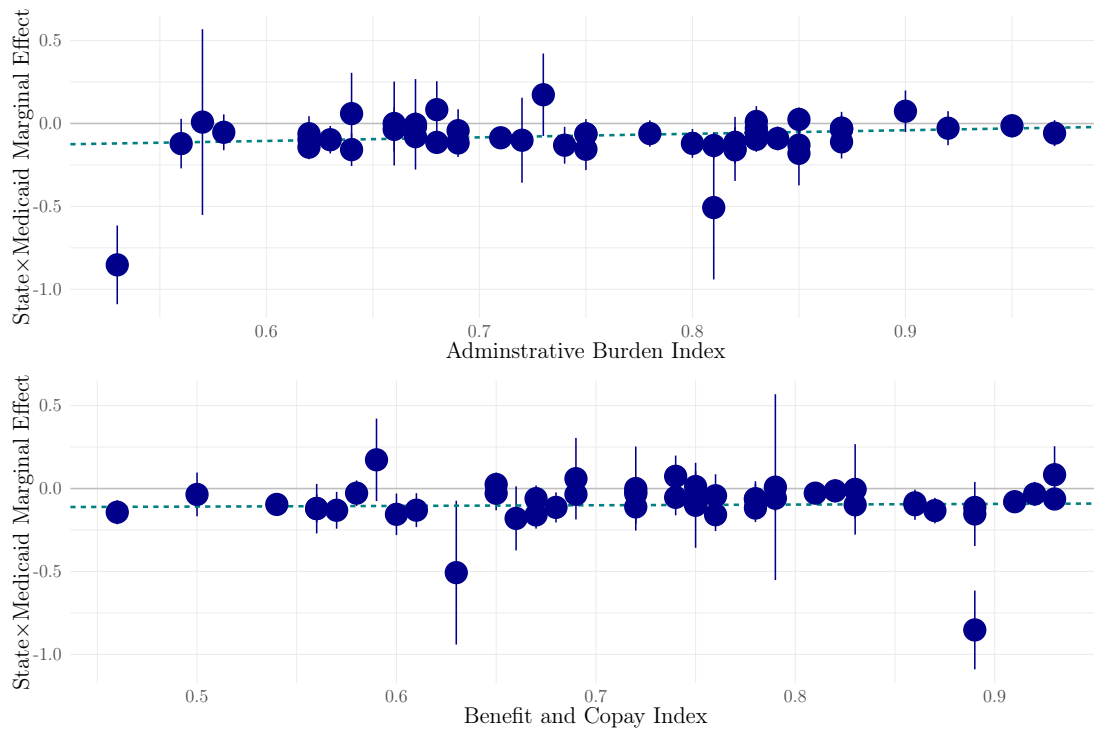
2SLS estimates above represent the results of estimating the naive LPM model described by equation (3.1), but using the percentage of Medicaid income eligibility cap in the first stage as an instrument to account for endogenous selection into Medicaid participation. Percent of Medicaid income eligibility cap is estimated as the ratio of the midpoint of the self-reported household income bin, and the state-of-residence's maximum Medicaid eligibility income cap based on the household dependents-conditional federal poverty level. Subset analysis follows the same column format as the primary results; covariates, fixed effects and standard errors clustering procedure are also identical to primary results.

3.4.2 Mechanisms

These initial results demonstrate a strong correlation between Medicaid status and the vaccination decision, but are also unclear and painting a considerably more complex picture than perhaps initially expected. The question remains for what mechanisms are actually driving these observations. If the particular administrative burden and costs of free care mechanisms are the primary drivers for the negative Medicaid coefficient as hypothesized, then I should be able to observe some degree of heterogeneity in the magnitude of this treatment effect across states and these variations in treatment should correlate to those states' respective mechanism indexes. To model this explicitly I re-estimate the entropy balanced model performed on the Medicaid and privately insured sub-sample but include a vector of State \times Medicaid interaction terms; I then plot these estimated interaction coefficients and their 95% confidence intervals in Figure 3.5 for each state against their average administrative burden index (top panel) and average benefit and copay index (bottom panel). To improve interpretability, the Rhode Island interaction term is omitted as the baseline reference point as it has both median administrative burden and benefit and copay scores. While it is difficult to discern any obvious positive relationships among these data, there does appear to be a *slight* upward trend in the state-level Medicaid marginal effects against the administrative burden index. At a minimum, I can observe that no State \times Medicaid interaction coefficients are significantly positive compared to the dozen that are significantly negative, suggesting Medicaid recipients are rendering vaccination decisions in a fairly consistent manner across the the country.

To estimate the magnitude of these mechanism effects formally, I again replicate the primary entropy balancing estimates on the Medicaid and uninsured sub-sample, but with the inclusion of AB \times Medicaid and BC \times Medicaid interaction terms to capture the marginal influence of administrative burden and benefits and copays on vaccination likelihood. Regression results are listed in Table 3.9. The results in column (4) are identical to the primary entropy balanced results from Table 3.3; column (5) is identical to (4) but includes the AB \times Medicaid and BC \times Medicaid interaction terms; column (6) is identical to (5), but

Figure 3.5: Estimated State×Medicaid coefficients against Administrative Burden and Benefits & Copay Indexes.



Coefficient estimates derived from an entropy balanced model on the Medicaid and privately insured populations sub-sample, with the inclusion of State×Medicaid interaction terms. The Rhode Island interaction term is omitted as the baseline reference point as it has both median administrative burden and benefit and copay scores. Heterogeneous, state-level estimates for the marginal influence Medicaid status on vaccination likelihood, relative to the Rhode Island median, are plotted against the Administrative Burden and Benefits & Copay Indexes from Fox, Feng, et al. (2020) for the respective states, along with their 95% confidence intervals.

substitutes the state-level per capita Medicaid expenditures (PC, in thousands of USD) for the administrative burden index and estimates the interaction term, $PC \times \text{Medicaid}$; column (8) includes all three interaction terms, $AB \times \text{Medicaid}$, $BC \times \text{Medicaid}$, and $PC \times \text{Medicaid}$. Interestingly, only the administrative burden index appears to be significant at the 5% level in describing state-by-state variation in vaccination propensities across all of the model specifications. Moreover, the estimated marginal effects for increased AB score are positive, which imply that reducing administrative burden barriers (a higher AB score) is associated with higher marginal likelihoods to vaccinate within the Medicaid population. I can observe an average increase in the likelihood to vaccinate when moving from the lowest (0.53 in Texas) to highest (0.97 in Illinois) index scores of 4.75 percentage points for administrative burden; or a 1.17 percentage point increase in vaccination propensities for a one standard deviation improvement in a state's administrative burden score.

In Table 3.12 I perform this same analysis, but on the Medicaid and uninsured populations sub-sample. Curiously, I observe the exact opposite outcomes, with the administrative burden score again being the only consistently significant factor influencing Medicaid vaccination outcome, but with relatively large coefficient estimates. One reason for these seemingly counterintuitive results could be that as a state's AB score improves, more of the uninsured population receives coverage and the selection into Medicaid effects become less pronounced. For instance, states with lower AB scores may have Medicaid populations which are more sickly (and more likely to vaccinate) than those that have less onerous enrollment procedures.

To make more sense of these mechanisms results, I perform some additional entropy balancing exercises following equation (3.2), but stratify by state-of-residence direct measures of Medicaid enrollment ease and list the estimates in Table 3.10. Enrollment ease data comes from Code for America (2019) and describe whether a given state permits joint application and processing for receipt of SNAP and Medicaid benefits, the number of screens required to click through to complete the online Medicaid application is at or below (above) the 25th (75th) percentile, and whether average time it takes to complete the online Medicaid application on a standard desktop computer is at or below (above) the

Table 3.9: Entropy balanced estimates for likelihood of receiving or intention to receive vaccine by Medicaid status and state-level Medicaid measures interactions (Medicaid recipients against privately insured).

	<i>Dependent variable:</i>			
	Vaccine Status			
	(4)	(5)	(6)	(7)
Medicaid	-0.069*** (0.007)	-0.129** (0.064)	-0.170*** (0.053)	-0.148** (0.072)
BC×Medicaid		-0.025 (0.052)		-0.024 (0.052)
AB×Medicaid		0.102* (0.055)	0.112** (0.055)	0.108* (0.056)
PC×Medicaid			0.002 (0.003)	0.002 (0.003)
Observations	79,647	79,647	79,647	79,647
R ²	0.22487	0.22500	0.22501	0.22501
Within R ²	0.08726	0.08742	0.08743	0.08743

Clustered (Week×State) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

((a)) All coefficient estimates are derived from a fixed effects linear probability model weighted according to the entropy balancing protocol outlined in equation (3.2). Analysis is performed on the entropy balanced sub-sample of Medicaid recipients and the privately-insured. BC and AB are the benefits and copay and administrative burden indexes collected from Fox, Feng, et al. (2020), PC is the 2019 per capita Medicaid expenditures by state; all three variables are interacted with the MEDICAID dummy in the regression. Covariates, fixed effects and standard errors clustering procedure are also identical to primary results.

25th (75th) percentile. When comparing states across these measures, only the differences between SNAP dual enrollment/processing and number of application screen measures are significant to the 5% level after performing a coefficient difference test. Somewhat surprisingly, states that permit greater ease when enrolling in either or both the SNAP and Medicaid programs demonstrate lower marginal vaccination rates among the Medicaid population. This could be a result of states with easier enrollment processes having a larger proportion of individuals who are historically underserved or have lower healthcare engagement levels. To more precisely gauge the effort required for enrollment, the "number of screens" variable distinguishes states in the highest or lowest quartile based on the actual volume of paperwork individuals must complete to enroll in Medicaid. I observed that on both the Medicaid vs uninsured, and Medicaid vs privately insured sub-samples, states which require the most amount of 'work' to apply for Medicaid benefits demonstrate marked declines in vaccination propensities in the Medicaid population. When stratified across application time quartiles, there do not appear to be any significant differences between states. These results generally appear to validate the crux of my argument: as states increase the bureaucratic burden that Medicaid recipients face, they become less inclined to utilize vaccination services.

3.5 Discussion & Conclusions

Overall my estimates appear to consistently demonstrate two common trends among the Medicaid population: After accounting for confounding factors and selection effects, Medicaid status appears to imbue the otherwise uninsured with a greater propensity to vaccinate; but that these benefits are trumped by program shortcomings when compared against conventional insurance enrollment. The results of the hesitancy motivation analysis bolster the argument that systemic mistrust is a driving factor in vaccine hesitancy. Taken together this suggests that the experience individuals accumulate and familiarity gained from accessing care through Medicaid do represent significant means of circumventing well-established informational barriers to vaccination (Dubé et al., 2013), but that these benefits are still insufficient compared to similar services offered by conventional insurance.

Table 3.10: State stratification entropy balancing analysis by Medicaid enrollment ease.

		<i>Dependent variable:</i>		
		Vaccine Status		
<i>Sub-sample:</i>		Medicaid Only & Uninsured	Pooled	Medicaid Only & Insured
Joint SNAP and Medicaid application/processing		0.034*** (0.013)	-0.042*** (0.010)	-0.044*** (0.011)
	Observations	12,043	29,079	24,996
Separate SNAP and Medicaid application/processing		0.044*** (0.016)	-0.024** (0.011)	-0.019 (0.013)
	Observations	6,191	12,705	10,739
# of screens to complete application (25th percentile)		0.053*** (0.014)	-0.024** (0.011)	-0.017 (0.011)
	Observations	8,453	17,991	15,444
# of screens to complete application (75th percentile)		0.027** (0.011)	-0.028*** (0.008)	-0.032*** (0.009)
	Observations	13,490	30,931	25,956
Time required to complete application (25th percentile)		0.042*** (0.010)	-0.028*** (0.007)	-0.026*** (0.008)
	Observations	8,453	17,991	15,444
Time required to complete application (75th percentile)		0.046*** (0.012)	-0.024** (0.010)	-0.024** (0.010)
	Observations	10,321	24,138	20,469

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

All coefficient estimates are derived from a fixed effects linear probability model weighted according to the entropy balancing protocol outlined in equation (3.2), stratified by state-of-residence measures of Medicaid enrollment ease. Enrollment ease data comes from Code for America (2019) and describe whether a given state permits joint application and processing for receipt of SNAP and Medicaid benefits, the number of screens required to click through to complete the online Medicaid application, and the average time it takes to complete the online Medicaid application on a standard desktop computer. Subset analysis follows the same column format as the primary results; covariates, fixed effects and standard errors clustering procedure are also identical to primary results.

I argue here that the comparative magnitude of the Medicaid-vs-uninsured and Medicaid-vs-insured causal estimates is evidence of the relative weight of the implicit costs (administrative burden, stigma, insurance discrimination by providers, etc.) associated with the receipt of Medicaid. Table 3.9 illustrates that, of the three mechanisms examined, only administrative burden appears to be a significant factor in capturing cross-state variations in vaccination rates. In both relative and absolute terms, variations in administrative burden appear to constitute approximately four times as much influence in the vaccination decision as benefits coverage and copays, which is also consistent with the zero-cost nature of the COVID-19 vaccine.

To estimate the economic magnitude of vaccine hesitancy demonstrated by Medicaid recipients, I perform some back of the envelope calculations using costs figures derived by Bartsch et al. (2021). Medicaid enrollment for May 2021 was 82.8 million³⁶; if I consider two counterfactual scenarios, one where Medicaid recipients were no more likely to vaccinate than the uninsured, and another where Medicaid recipients were no less likely than the insured to vaccinate, then my estimates would imply there have been 2.815 million more, and 5.713 million fewer vaccinations respectively (or 0.85% and 1.72% of the total US population) based on the entropy balanced figures from Table 3.3. Using the most conservative figures for averages in moving from 50% to 70% vaccine coverage (Bartsch et al., 2021), I find this corresponds to 390,200 fewer cases and 1,662 fewer deaths than if Medicaid recipients were no more likely to vaccinate than the uninsured; and 791,800 more cases and 3,373 more deaths than if Medicaid recipients were no less likely to vaccinate than the insured over the observed January to September, 2021 time frame. With respect to reducing administrative burden, if all states improved their enrollment procedures so as to increase their AB scores by one standard deviation, this would correspond to 971,900 more vaccinations, 134,700 fewer COVID-19 cases, and 574 fewer deaths. The increased proclivity to vaccinate that Medicaid status affords clearly saves lives, but could stand to save substantially more if it were as persuasive as conventional insurance.

³⁶<https://www.kff.org/coronavirus-covid-19/issue-brief/analysis-of-recent-national-trends-in-medicaid-and-chip-enrollment/>

This analysis warrants consideration of several potential limitations and contextual factors. Medicaid eligibility criteria extend beyond mere income and household size, hinting at possible missing variables in the balancing models that could restrict the causal interpretation of the findings. To address this, I utilize the Medicaid income cap percentage as a tool for introducing exogenous variation in enrollment, applying it within both regression discontinuity and instrumental variable frameworks. Despite the challenges posed by the imprecision of self-reported income, which may affect the regression discontinuity approach's effectiveness, the instrumental variable analysis consistently reveals lower vaccination rates among Medicaid enrollees in both pooled and insured groups after accounting for selection effects. Furthermore, changes in Medicaid enrollment policies during the pandemic, including reduced churn and relaxed eligibility criteria, might be influencing the estimates. Lastly, the study does not account for how shifts in administrative burdens over time could affect the results, adding another layer of complexity to their interpretation.

The policy implications of my results here highlight the urgency in reducing or eliminating administrative burden in the Medicaid enrollment and care utilization processes. From a public health perspective, the magnitude of the costs of vaccine initiative failure is substantial; but to portray hesitancy as irrational appears to be unsubstantiated. As with all forms of care utilization, individuals form expectations on care quality and costs based on experience and institutional familiarity. Even if Medicaid recipients possess greater institutional familiarity, this alone does not (at least not entirely) offset the general unpleasantness of utilizing Medicaid care. Some form of welfare analysis that includes, for instance, the administrative burden measurements could be helpful in identifying the marginal willingness-to-pay for vaccination among the Medicaid population relative to the increased time and effort costs of receiving care. As well, including a formal measure of insurance discrimination faced by Medicaid recipients from care providers could further distinguish between the relative size of implicit administrative and utilization costs.

References

- Alcalá, H. E., & Cook, D. M. (2018). Racial discrimination in health care and utilization of health care: a cross-sectional study of california adults. *Journal of general internal medicine*, *33*(10), 1760–1767.
- Alcalá, H. E., Ng, A. E., Gayen, S., & Ortega, A. N. (2020). Insurance types, usual sources of health care, and perceived discrimination. *The Journal of the American Board of Family Medicine*, *33*(4), 580–591.
- Alharbi, A. G., Khan, M. M., Horner, R., Brandt, H., & Chapman, C. (2019). Impact of medicaid coverage expansion under the affordable care act on mammography and pap tests utilization among low-income women. *PloS one*, *14*(4), e0214886.
- Allen, E. M., Call, K. T., Beebe, T. J., McAlpine, D. D., & Johnson, P. J. (2017). Barriers to care and healthcare utilization among the publicly insured. *Medical care*, *55*(3), 207.
- Allen, H., Gordon, S. H., Lee, D., Bhanja, A., & Sommers, B. D. (2021). Comparison of utilization, costs, and quality of medicaid vs subsidized private health insurance for low-income adults. *JAMA network open*, *4*(1), e2032669–e2032669.
- Antonisse, L., Garfield, R., Rudowitz, R., & Artiga, S. (2019). The effects of medicaid expansion under the aca: updated findings from a literature review. *KFF Issue Brief*.
- Aslim, E. G., Fu, W., Liu, C.-L., & Tekin, E. (2022). *Vaccination policy, delayed care, and health expenditures* (Tech. Rep.). National Bureau of Economic Research.
- Aslim, E. G., Fu, W., Tekin, E., & You, S. (2023). *From syringes to dishes: improving food security through vaccination* (Tech. Rep.). National Bureau of Economic Research.
- Baicker, K., Finkelstein, A., et al. (2019). The impact of medicaid expansion on voter participation: Evidence from the oregon health insurance experiment. *Quarterly Journal of Political Science*, *14*(4), 383–400.
- Bartsch, S. M., Wedlock, P. T., O’Shea, K. J., Cox, S. N., Strych, U., Nuzzo, J. B., ... others (2021). Lives and costs saved by expanding and expediting covid-19

- vaccination. *The Journal of Infectious Diseases*.
- Calonico, S., Cattaneo, M. D., & Farrell, M. H. (2020). Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs. *The Econometrics Journal*, 23(2), 192–210.
- Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6), 2295–2326.
- Calonico, S., Cattaneo, M. D., & Titiunik, R. (2015). Optimal data-driven regression discontinuity plots. *Journal of the American Statistical Association*, 110(512), 1753–1769.
- Casey, R., & Wyden, R. (2022, March). *Letter to the gao requesting study on medicaid covid-19 vaccination rates*. Retrieved from <https://khn.org/wp-content/uploads/sites/2/2022/03/2022.03.25-Casey-Wyden-GAO-Request-Letter-on-Medicaid-COVID-Vax-Rates.pdf> (Letter to the Comptroller General of the United States, Government Accountability Office)
- Churchill, B. F. (2021). Insurance coverage, provider contact, and take-up of the hpv vaccine. *American Journal of Health Economics*, 7(2), 222–247.
- Code for America. (2019). *Bringing social safety net benefits online*. Retrieved from <https://codeforamerica.org/explore/bringing-social-safety-net-benefits-online/>
- Cook, R. L., Zhang, J., Mullins, J., Kauf, T., Brumback, B., Steingraber, H., & Mallison, C. (2010). Factors associated with initiation and completion of human papillomavirus vaccine series among young women enrolled in medicaid. *Journal of Adolescent Health*, 47(6), 596–599.
- Dubé, E., Laberge, C., Guay, M., Bramadat, P., Roy, R., & Bettinger, J. A. (2013). Vaccine hesitancy: an overview. *Human vaccines & immunotherapeutics*, 9(8), 1763–1773.
- Fox, A. M., Feng, W., Zeitlin, J., & Howell, E. A. (2020). Trends in state medicaid eligibility, enrollment rules, and benefits: Study examines state-level trends in medicaid eligibility, benefits, and administrative burden, both before and after

- implementation of the affordable care act. *Health Affairs*, 39(11), 1909–1916.
- Fox, A. M., Stazyk, E. C., & Feng, W. (2020). Administrative easing: Rule reduction and medicaid enrollment. *Public Administration Review*, 80(1), 104–117.
- Fridman, A., Gershon, R., & Gneezy, A. (2021). Covid-19 and vaccine hesitancy: A longitudinal study. *PloS one*, 16(4), e0250123.
- Galewitz, P. (2022). *From alabama to utah, efforts to vaccinate medicaid enrollees against covid run into obstacles*. Kaiser Family Foundation Health News. Retrieved from <https://kffhealthnews.org/news/article/medicaid-covid-vaccine-obstacles-states/>
- Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political analysis*, 20(1), 25–46.
- Han, X., Call, K. T., Pintor, J. K., Alarcon-Espinoza, G., & Simon, A. B. (2015). Reports of insurance-based discrimination in health care and its association with access to care. *American journal of public health*, 105(S3), S517–S525.
- Hill, H. A., Elam-Evans, L. D., Yankey, D., Singleton, J. A., & Kang, Y. (2017). Vaccination coverage among children aged 19–35 months—united states, 2016. *Morbidity and Mortality Weekly Report*, 66(43), 1171.
- Hill, H. A., Elam-Evans, L. D., Yankey, D., Singleton, J. A., & Kang, Y. (2018). Vaccination coverage among children aged 19–35 months—united states, 2017. *Morbidity and Mortality Weekly Report*, 67(40), 1123.
- Hollingsworth, A., Soni, A., Carroll, A. E., Cawley, J., & Simon, K. (2019). Gains in health insurance coverage explain variation in democratic vote share in the 2008-2016 presidential elections. *PloS one*, 14(4), e0214206.
- Hornsey, M. J., Finlayson, M., Chatwood, G., & Begeny, C. T. (2020). Donald trump and vaccination: The effect of political identity, conspiracist ideation and presidential tweets on vaccine hesitancy. *Journal of Experimental Social Psychology*, 88, 103947.
- Khan, M. S., Ali, S. A. M., Adelaine, A., & Karan, A. (2021). Rethinking vaccine hesitancy among minority groups. *The Lancet*, 397(10288), 1863–1865.

- Mazurenko, O., Balio, C. P., Agarwal, R., Carroll, A. E., & Menachemi, N. (2018). The effects of medicaid expansion under the aca: a systematic review. *Health Affairs*, *37*(6), 944–950.
- Moynihan, D. P., Herd, P., & Ribgy, E. (2016). Policymaking by other means: Do states use administrative barriers to limit access to medicaid? *Administration & Society*, *48*(4), 497–524.
- Naderalvojud, B., Shah, N. D., Mutanga, J. N., Belov, A., Staiger, R., Chen, J. H., . . . Hernandez-Boussard, T. (2023). Trends in influenza vaccination rates among a medicaid population from 2016 to 2021. *Vaccines*, *11*(11), 1712.
- Nguyen, T.-L., Collins, G. S., Spence, J., Daurès, J.-P., Devereaux, P., Landais, P., & Le Manach, Y. (2017). Double-adjustment in propensity score matching analysis: choosing a threshold for considering residual imbalance. *BMC medical research methodology*, *17*(1), 1–8.
- O'Malley, C. D., Shema, S. J., Clarke, L. S., Clarke, C. A., & Perkins, C. I. (2006). Medicaid status and stage at diagnosis of cervical cancer. *American journal of public health*, *96*(12), 2179–2185.
- Soares, P., Rocha, J. V., Moniz, M., Gama, A., Laires, P. A., Pedro, A. R., . . . Nunes, C. (2021). Factors associated with covid-19 vaccine hesitancy. *Vaccines*, *9*(3), 300.
- Stoecker, C., Stewart, A. M., & Lindley, M. C. (2017). The cost of cost-sharing: the impact of medicaid benefit design on influenza vaccination uptake. *Vaccines*, *5*(1), 8.
- Stuber, J., & Kronebusch, K. (2004). Stigma and other determinants of participation in tanf and medicaid. *Journal of Policy Analysis and Management*, *23*(3), 509–530.
- Thorburn, S., & De Marco, M. (2010). Insurance-based discrimination during prenatal care, labor, and delivery: perceptions of oregon mothers. *Maternal and child health journal*, *14*(6), 875–885.
- Weech-Maldonado, R., Hall, A., Bryant, T., Jenkins, K. A., & Elliott, M. N. (2012). The relationship between perceived discrimination and patient experiences with health care. *Medical care*, *50*(9 0 2), S62.

-
- Wright, B. J., Carlson, M. J., Edlund, T., DeVoe, J., Gallia, C., & Smith, J. (2005). The impact of increased cost sharing on medicaid enrollees. *Health Affairs*, *24*(4), 1106–1116.
- Zhao, Q., & Percival, D. (2017). *Entropy balancing is doubly robust* (Vol. 5) (No. 1). De Gruyter.

3.7 Appendices

3.7.1 Propensity score matched model results

As an alternative to the entropy balancing approach used in the primary component of my analysis, I additionally employ propensity score matching (PSM) to account for any endogeneity in the MEDICAID_i treatment variable.

I estimate the probability of being ‘treated’ (enrolled in Medicaid) given a set of pre-treatment covariates according to the following logistic regression:

$$\text{MEDICAID}_i = X_i\delta_1 + \text{STATE}_i\gamma + \epsilon_{1i} \quad (3.4)$$

where, as with (3.1), X_i corresponds to vectors of individual and household demographic variables, and STATE is a vector of state-of-residency indicators to ensure matching is performed to locate within state controls.

To eliminate potential biasing from selection into treatment, I use the estimated propensity scores from (3.4) for each observation and employ a nearest neighbor matching algorithm to pair each ‘treated’ individual with their most demographically similar untreated control. Then to ensure robustness and account for potential imbalance within the one-to-one matched sub-sample (Nguyen et al., 2017), I run the following LPM and fixed effects logit:

$$y_{i,t} = \alpha_2 + \beta \text{MEDICAID}_i + X_i\delta_2 + \phi_{s,t} + \epsilon_{2i,t} \quad (3.5)$$

which follows a similar variable specification as (3.1). The results of these regressions are listed in Table 3.11. To summarize: These results are virtually identical to the primary entropy balancing estimates, with the same significance levels, signs, and magnitudes for the coefficients. The Medicaid population exhibits a 3.3 percentage point higher likelihood to vaccinate compared to the matched uninsured population, and a 6.5 percentage point lower likelihood to vaccinate compared to the matched insured population. When examining the net effects on the matched pooled population (row 2), these PSM results are slightly

larger than the entropy balancing estimates, indicating a 3.6 percentage point decline in vaccination likelihoods compared to broader, Medicaid-similar population.

Table 3.11: Propensity score matched estimates for likelihood of receiving or intention to receive vaccine by insurance coverage type.

		<i>Dependent variable:</i>		
		Vaccine Status		
<i>Sub-sample:</i>		Medicaid Only & Uninsured	Pooled	Medicaid Only & Insured
<i>(a) Linear probability estimates</i>	Medicaid	0.029*** (0.009)	-0.037*** (0.004)	-0.065*** (0.004)
	Observations	28,352	56,482	56,482
	R ²	0.18311	0.18618	0.18971
	Within R ²	0.07009	0.08671	0.09351
<i>(b) Fixed effects logistic regression estimates</i>	Logit coefficients	0.168*** (0.048)	-0.197*** (0.022)	-0.371*** (0.024)
	Estimated Marginal effects	0.033*** (0.009)	-0.036*** (0.005)	-0.065*** (0.006)
	Observations	28,208	56,372	56,393
	Squared Correlation	0.18123	0.18833	0.19345
	Pseudo R ²	0.14537	0.15648	0.16301
	BIC	42,778.3	72,058.2	70,157.5

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

((a)) All coefficient estimates are derived from a fixed effects linear probability model estimated on propensity score matched samples produced according to equation (3.4). Estimates in column one describe the marginal change in vaccination likelihood of Medicaid status compared to the uninsured population; column three estimates have the same interpretation, but compare vaccination outcomes between the Medicaid and privately-insured populations; column two estimates these changes for the pooled uninsured, Medicaid, and privately-insured populations. Covariates include: individual age, sex, marriage status, ethnicity, race, logged household income, home ownership status, household size, number of children in household, employment status, whether they had contracted COVID at some point, and whether they had difficulty with paying bills. Estimates include state-week and MSA-week fixed effects; all standard errors are clustered at the state-week level.

((b)) All coefficient estimates are derived from a fixed effects logistic regression model estimated on propensity score matched samples produced according to equation (3.4). The estimated marginal effects in the second row have the same interpretation as the primary results outlined in panel (a).

Table 3.12: Entropy balanced estimates for likelihood of receiving or intention to receive vaccine by Medicaid status and state-level Medicaid measures interactions (Medicaid recipients against uninsured).

	<i>Dependent variable:</i>			
	Vaccine Status			
	(4)	(5)	(6)	(7)
Medicaid	0.034*** (0.007)	0.072 (0.077)	0.171*** (0.066)	0.088 (0.090)
BC×Medicaid		0.096 (0.064)		0.096 (0.064)
AB×Medicaid		-0.143** (0.066)	-0.162** (0.066)	-0.148** (0.068)
PC×Medicaid			-0.001 (0.004)	-0.001 (0.004)
Observations	42,417	42,417	42,417	42,417
R ²	0.19220	0.19255	0.19247	0.19256
Within R ²	0.08056	0.08096	0.08086	0.08097

Clustered (Week×State) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

All coefficient estimates are derived from a fixed effects linear probability model weighted according to the entropy balancing protocol outlined in equation (3.2). Analysis is performed on the entropy balanced sub-sample of Medicaid recipients and the uninsured. BC and AB are the benefits and copay and administrative burden indexes collected from Fox, Feng, et al. (2020), PC is the 2019 per capita Medicaid expenditures by state; all three variables are interacted with the MEDICAID dummy in the regression. Subset analysis follows the same column format as the primary results; covariates, fixed effects and standard errors clustering procedure are also identical to primary results.