Designing Pandemic Vaccination: Private and Public Returns in Optimal Allocation

Rishabh Subramanian*[1]

[1]University of Chicago
*rishabhs@uchicago.edu

August 2025

Abstract

Allocating scarce vaccines in a pandemic requires policymakers to resolve a fundamental trade-off: whether to protect the vulnerable from severe illness and death, or to curb transmissions by targeting socially active groups. This paper provides the population-level evidence needed to quantify both margins for the elderly (65+) and younger adults (18–65). Using monthly data for all U.S. counties from January 2020 to July 2022, I exploit the staggered rollout of vaccines and the differential timing across age groups, combining a triple-difference estimation with difference-in-differences estimates for each age group. A 10% rise in vaccination reduces deaths, ICU admissions, and hospitalizations by 3-, 13-, and 8-fold more for older than younger adults, respectively, while infections fall 3-fold more among younger adults. All these effects are measured one year after roll-out, underscoring the importance of long-run impacts for policy design. *Counterfactual simulations* based on these causal estimates reveal the following: prioritizing the elderly minimizes severe outcomes, targeting the young curbs spread, and proportional allocation delivers no clear win. By pairing credible long-run estimates of the inter-generational trade-off with forward-looking simulations, this study provides an actionable, evidence-based blueprint for optimal vaccine allocation in future pandemics.

JEL: I12, I18, J14
Keywords: Vaccine, Death, Elderly, Pandemic, COVID-19, Externalities, Hospital, ICU, Public Good, Private Benefits, Trade-off

*I would like to thank Professors Tarun Jain, Victor Lima, Sriram Venkataraman, and Luigi Zingales for their valuable comments and feedback. All errors are mine.

1 Introduction

Consider a policymaker confronting the next devastating pandemic. With the COVID-19 pandemic causing over a million deaths and an estimated $14 trillion in economic losses in the United States (US) alone,[1] the stakes for designing sound health policy will be immense. At the same time, public trust in government policies declined sharply during the COVID-19 pandemic,[2] leaving future policymakers with the challenge of preventing catastrophic losses while sustaining public confidence. This challenge is not distant or hypothetical. The World Health Organization (WHO) has warned that the world must prepare for outbreaks that could be even deadlier than COVID-19.[3] Moreover, scientists point out that ongoing deforestation and biodiversity loss are increasing the likelihood of zoonotic spillovers and global contagion events.[4] Maintaining trust and preventing the social and economic toll in future pandemics will require policy grounded in rigorous evidence.

Vaccination is widely regarded as a cornerstone of pandemic defense.[5] Yet, as vaccine supply is often limited—especially at the onset of a pandemic—policymakers confront a fundamental economic dilemma. *Vaccines serve as both private protective goods, offering direct protection to recipients, and public goods replete with externalities.* For older individuals, who face the highest health risks, vaccination yields large private returns by preventing severe illness and death. For younger adults, who face lower individual risks but have higher contact rates, vaccination can generate sizable public-good benefits by interrupting chains of transmission and increasing mobility (Mukherjee et al., 2021). Slowing viral spread also carries major macroeconomic consequences, bolstering labor supply, production, output, employment, and price stability (Adda, 2016). The central question is then as follows: should scarce vaccines be directed toward the elderly to minimize severe outcomes, or toward the young to curb spread as well as limit economic damage?

Biological differences in vaccine responsiveness exacerbate this dilemma. Older adults often exhibit diminished vaccine responsiveness due to *immunosenescence*, the age-related decline in the functioning of the immune system. In comparison, younger adults may gain robust immunity from vaccination and, through higher rates of social contact, contribute disproportionately more to herd protection. These contrasts mean that both biological and behavioral factors must be weighed when deciding how to allocate scarce vaccines under evolving epidemiological risk.

While numerous countries, guided by WHO recommendations,[6] prioritized the elderly for COVID-19 vaccination, was that policy optimal? Current evidence offers policymakers limited guidance on how to resolve the inter-generational trade-off. *Resolving this dilemma requires evidence drawn from population-level analysis* because only such data can jointly capture the private benefits to high-risk groups and the spillovers from reduced transmission in lower-risk groups.

[1] See Walmsley et al. (2023) and https://covid.cdc.gov/covid-data-tracker

[2] See Aassve et al. (2024); Perlis et al. (2024)

[3] See https://www.independent.co.uk/news/health/who-pandemic-warning-covid-b2344635.html

[4] Tollefson (2020) explains that deforestation and extinctions are likely to increase the occurrence of pandemics.

[5] See García-Montero et al. (2021); Machado et al. (2022)

[6] See https://www.who.int/news-room/questions-and-answers/item/coronavirus-disease-(covid-19)-vaccine-access-and-allocation

Without this dual quantification, policymakers cannot credibly compare allocation strategies or weigh severe outcomes against community spread. Most existing studies rely on clinical or survey-based data that, by definition, cover only a subset of the population and cannot fully capture the population-level effects.

This paper provides the necessary evidence to fill these gaps. Using monthly, population-level data from over 3,000 U.S. counties from January 2020 to July 2022, this study combines a triple-difference design using age-based prioritization in vaccine roll-out with difference-in-difference estimates, which utilize variation in vaccine uptake across counties and time.

Each outcome examined in this paper reflects a distinct dimension of vaccine impact. Deaths, ICU admissions, and hospitalizations represent the private benefit of vaccination, shielding individuals from severe or fatal illness. In contrast, infections (cases) capture the public-good component of vaccination, as preventing infections in one individual can reduce the transmission risk for others. As a growing literature underscores the significance of transmission-related externalities in pandemic policy, studying this public externality as the flip side of the policy trade-off is crucial.[7] Using county-level analysis and confirming that all the results replicate when the tests are conducted at the MSA level, this study accounts for transmission-related externalities with geographic precision, offering causal evidence of such effects.

The results reveal a striking asymmetry: vaccination provides consistently greater reductions in deaths and severe illness among the elderly, while infections decline significantly more among the younger adults—reflecting their central role in transmission and their stronger immune response to vaccination. This contrast confirms the policymaker's dilemma of whether to prioritize reducing severity or curbing spread, and highlights that the answer depends not only on vaccine supply, but also on which outcome is most valued at a given stage of the pandemic.

My empirical strategy is designed to deliver credible causal estimates of the long-run health effects of vaccination. The core specification is a triple-difference framework that compares the impact of vaccination on outcomes between older and younger adults within the same county and month. This ensures that identification comes from within-county, cross-age differences in vaccination intensity, under shared local conditions. To assess the robustness of this approach, I estimate the triple-difference with alternative fixed effects structures—including (county, time), (county, age group) as well as state-specific time trends added to county and time fixed effects. The fixed effects for each (county, time) pair flexibly absorb all county-specific shocks in each time period, such as local virus waves, health system strain, or behavioral shifts that affect both age groups equally. The other two specifications account for the possibility that age-specific behavioral responses or policy shocks might bias the estimates. The key results remain consistent across these alternative specifications, reinforcing the credibility of the triple-difference design and indicating that the findings are not sensitive to a particular fixed-effects model.

While the triple-difference strategy provides the most rigorous identification, the accompanying difference-in-differences estimates also appear to yield credible causal effects. Four factors support

[7] See Cakmakli et al. (2021); Andersson et al. (2021); Bundorf et al. (2022)

this interpretation. First, the DiD estimates display a sharp reversal in the *slope* of health outcomes at the onset of vaccination, a pattern consistent with a causal response rather than a continuation of prior trends. Second, most policy variation during the pandemic occurred at the state level, while the identifying variation in vaccine uptake arises at the county level. This county-level variation is plausibly quasi-random: as shown by Allcott et al. (2024), neither policy mandates nor rates of voluntary social distancing explain much of the geographic variation in pandemic outcomes, which is instead driven by exogenous county characteristics. Since these characteristics are absorbed by county fixed effects, the DiD estimates are unlikely to be confounded by unobserved local policy or behavior. Third, examination of the dynamic effects reveals that while the effect for infections, hospitalizations, and ICU admissions starts manifesting within a month after vaccination, the effect for deaths only starts manifesting two months after vaccination. This *staggered onset* strengthens the interpretation that the observed effects reflect the natural epidemiological unfolding of vaccine impact. Finally, the DiD estimates remain similar to those derived using the more rigorous triple-difference estimates.

By combining triple-difference and difference-in-differences designs, the paper isolates both the relative marginal returns across groups and the absolute gains within each group, over a one-year horizon. Deaths, ICU admissions, and hospitalizations decrease by 3-fold, 13-fold, and 8-fold more for the elderly than the younger adults, even one year after a 10% increase in vaccination, highlighting the high marginal returns from targeting vulnerable groups. Conversely, infections fall three times as much among younger adults, reflecting their role in spreading disease. These effects are not only large in magnitude but persist for a year, offering rare long-run evidence that is especially valuable in light of concerns about waning immunity (Dolgin et al., 2021). Indeed, the effects after one year are more robust and consistent than the effects I obtain over shorter horizons.

To translate these findings into actionable insights, I *simulate counterfactual scenarios* under three vaccine allocation strategies: prioritizing only the elderly, prioritizing only younger adults, and distributing doses proportional to population shares. These counterfactual scenarios use actual vaccine supply by county and month, and apply the estimated treatment effects to county-level populations. The results show that prioritizing the elderly yields the greatest reduction in mortality and critical care burden, while prioritizing the young results in the largest decline in infections but at the cost of significantly higher deaths. A proportional allocation achieves moderate improvements in both outcomes but does not dominate on either front. These findings underscore the importance of age-based targeting during the early phases of vaccine roll-out and illustrate the practical consequences of different policy choices in the face of constrained supply and distinct age-group vulnerabilities.

This paper provides population-wide causal evidence that jointly quantifies the private, protective returns to vaccination and the public-good benefits from reduced transmission. By linking these heterogeneous effects across age groups to counterfactual policy simulations, this study moves beyond measurement to offer directly actionable guidance for vaccine allocation under scarcity. Whereas prior work has emphasized short-term outcomes, relied on clinical or survey data that

cannot capture population-level spillovers, and offered little evidence across demographic groups, this paper provides long-run estimates that reveal how benefits diverge sharply between the elderly and the young. In doing so, the paper connects reduced-form evidence to normative policy design, offering the first empirical foundation for optimal vaccination strategies in pandemics.

This paper is organized as follows. The next section provides an overview of the literature, specifically highlighting this study's contribution to the literature. Section 3 describes the vaccine rollout in the U.S. as it forms the basis for my empirical strategy. Section 4 describes the data and the proxies, while section 5 explains the empirical strategy. Section 6 reports the main results, while section 7 discusses additional robustness tests. Section 8 estimates the economic magnitudes, while section 9 describes the counter-factual scenarios I estimate. Since this study focuses on the U.S., section 10 provides a discussion about how to apply the lessons from this study to other country contexts. Section 11 concludes.

2 Review of Literature

A central challenge in pandemic response is the allocation of scarce vaccines across populations that differ both in vulnerability to severe illness and in their role in disease transmission. Theoretical and simulation-based studies have examined optimal vaccine allocation under scarcity, highlighting the potential tension between private protection and public externalities (Acemoglu et al., 2021; Alvarez et al., 2021; Cakmakli et al., 2021; Chang et al., 2021; Chetty et al., 2024). These contributions underscore the importance of targeting but rely on stylized assumptions and calibrated models, leaving open the question of how such trade-offs manifest in real-world populations. To bridge this gap, this study builds on a broader economics literature that has examined the balance of private and public returns in health policy (Feldstein, 1973; Manning et al., 1987; Feldman and Dowd, 1991; Cutler and Gruber, 1996; Eichner, 1998; Brown and Finkelstein, 2008; Card et al., 2008; Almond et al., 2010; Shi, 2024). Foundational work in this tradition shows how health policies can simultaneously deliver private benefits to individuals while generating broader spillovers, a perspective I extend by quantifying both protective and transmission-reducing returns to vaccination.

Recent empirical work has examined the broader health and economic impacts of the COVID-19 pandemic and vaccination campaigns (Chetty et al., 2024). Two studies are particularly relevant: Davies and Zhou (2023) analyze demographic variation in vaccine effects but not policy-driven allocation, while Kim and Lee (2022) focus on cross-country strategies without exploiting within-country heterogeneity. Related work highlights broader behavioral and economic channels—for example, Mukherjee et al. (2021) show that vaccination increased mobility, particularly among younger cohorts. Taken together, this empirical literature sheds light on important consequences of vaccination campaigns but does not quantify the allocation trade-offs policymakers face under scarcity.

Finally, this study connects to broader work on pandemic behavior, health care utilization,

and vaccine response (Okeke, 2022; Kim and Lee, 2022). While these contributions provide valuable context on vaccine coverage and behavioral heterogeneity (Bundorf et al., 2022; Hortaçsu et al., 2021; Andersson et al., 2021), they do not estimate differential returns to vaccination across population groups. By bridging insights from theory, health economics, and empirical evidence on pandemic behavior, this paper demonstrates how targeting distinct populations yields sharply divergent health returns—evidence directly relevant for the design of optimal allocation policies.

3 COVID-19 Vaccine Roll-out in U.S. Counties

Understanding the roll-out of COVID-19 vaccines across U.S. counties—and how it varied between age groups—is essential to understand the variation that is employed for identification. This section describes the same.

3.1 Age-based Prioritization in Vaccination

The U.S. COVID-19 vaccination campaign began in December 2020, with early phases explicitly prioritizing older adults (65+) and high-risk populations. State and county health departments implemented these priorities within local supply constraints, resulting in substantial differences in roll-out timing and speed across the country. As vaccine supply expanded, eligibility broadened to younger adults (18–65), but uptake continued to be shaped by local demographics, political preferences, and institutional capacity. These institutional features created systematic differences in vaccination rates between age groups within the same county and month—the variation central to this paper's identification strategy. Figures 1 and 2 document key features of this variation: both over time and across space, and at both the intensive and extensive margins.

These figures are based on census data on vaccinations obtained from CDC. This data is available daily at the county level. As COVID-19 vaccination started in the U.S. in Dec-2020, the vaccination data is also available from this month. This data provides the number and percentage of individuals in the county who have received one dose, two doses, and the booster dose. This data is provided for the following two age groups: 18+ years and 65+ years.

3.2 Temporal Trends in Vaccine Uptake

Figure 1 presents monthly trends in COVID-19 vaccine uptake for the elderly (65+ years) and the younger adults (18-64 years) across U.S. counties from Jan-2020 through Jul-2022. As the vaccination for both the elderly (65+ years) and the younger adults plateaued by early 2022, I use data on vaccination until Jul-2022. The figure is divided into three panels showing the share of each group that received: (i) two doses (top panel), (ii) at least one dose (middle panel), and (iii) a booster dose (bottom panel).

In all three panels, elderly individuals consistently exhibit faster and higher uptake of vaccines compared to the 18–65 age group, reflecting the age-based prioritization embedded in national and state-level vaccine roll-out strategies. The top panel shows the cumulative share of each population

that completed the primary two-dose regimen. Vaccination among the elderly accelerated rapidly after Jan-2021, crossing 50% coverage by mid-2021 and plateauing near 70% by early 2022. In contrast, the 18–65 group reached only about 55% two-dose coverage over the same period.

The middle panel, tracking receipt of at least one dose, shows similar patterns but highlights an earlier and steeper rise for the elderly beginning in the first quarter of 2021. This suggests that the elderly were not only prioritized earlier, but also responded with higher initial willingness to vaccinate. The gap between groups persisted through early 2022 before both curves leveled off, consistent with saturation effects.

The bottom panel shows uptake of booster doses, which became widely available in the second half of 2021. Here, the divergence between age groups is stark: booster uptake among the elderly climbs steadily to around 50% by mid-2022, while booster coverage for the 18–65 group remains below 25%. This widening gap reinforces the age-targeted structure of vaccination between the groups.

These trends are central to my identification strategy. The earlier and higher vaccine uptake among the elderly provides meaningful variation in vaccination rates across age groups within a (county, month)—variation that this study exploits to estimate heterogeneous treatment effects using both triple-difference and difference-in-differences designs.

3.3 Cross-Sectional and Temporal Variation Relevant for Identification

Figure 2 complements the insights obtained from these trends by visualizing two key parts of the variation in two-dose vaccine coverage across counties: (i) the vaccine uptake across U.S. counties, and (ii) the cross-sectional heterogeneity within each month, both disaggregated by age group. The top panel shows the number of counties where at least 50% of the relevant age group had received two doses of the COVID-19 vaccine. The red line (diamonds) represents counties reaching this threshold for the elderly (age 65+), while the blue line (circles) tracks the same metric for younger adults (age 18–65). We observe here as well that elderly vaccination surged early in 2021, with over 2,000 counties surpassing 50% coverage by Mar-2021, following prioritized roll-out policies. In contrast, uptake among the 18–65 group was slower, with only a few hundred counties crossing the 50% threshold by mid-2021, and many catching up only by early 2022. By Jul-2022, nearly all counties had surpassed this threshold for the elderly, while younger adults lagged in several counties. This divergence in early vaccination intensity across counties and age groups creates rich identifying variation that feeds our empirical designs

The bottom panel of Figure 2 presents monthly box plots of county-level two-dose coverage from Jan-2021 to Jul-2022, again separately for the elderly (pink boxes) and the younger adults (blue boxes). These box plots capture the cross-sectional variation in vaccination rates across counties at each point in time. Several important patterns emerge. First, median two-dose coverage among the elderly (the bold line in each pink box) consistently exceeds that of the younger group throughout the sample period. Second, the inter-quartile range for the elderly is initially wide but narrows as coverage saturates, while the 18–65 group exhibits sustained dispersion in coverage even into late

2022. Third, the top whiskers (90th percentile) for both groups approach or exceed 80% coverage by mid-2022, but the bottom quartile for younger adults remains well below 40% in many months. These persistent differences in distribution reflect local demographic, political, and institutional factors influencing vaccine uptake.

Taken together, these figures establish both the institutional basis for focusing on age-based heterogeneity and the empirical conditions necessary for causal identification.

4 Data and Proxies

I construct the census data on health outcomes from the case surveillance data on COVID-19 put out by the Centers for Disease Control and Prevention (CDC) for public use. This *anonymized* data is available at the individual level and captures whether the individual has a confirmed case of COVID-19 or not, whether the individual is symptomatic or not, whether the individual had to be hospitalized or admitted to the Intensive Care Unit (ICU) of a hospital, the date when the case was reported, the age group of the individual (0-17 years, 18-49 years, 50-64 years, and 65+ years), and the county and state the individual is located in. As mentioned above, the census data on vaccinations is also obtained from CDC. Crucially, the use of census-based (population-level) data *significantly reduces concerns about sample selection bias* in this study.[8]

4.1 Balanced panel at the (county, month, age group) level

I aggregate the case surveillance and vaccination data to the level of a (county, month). Due to individual-level confounding factors, the econometric concerns involved in identifying the causal effects of COVID-19 vaccine on health outcomes are pronounced when the analysis is done at the level of an individual (see Lipsitch and Dean (2020) and Dean et al. (2021) for details). When the data is aggregated to the level of a county, such individual-level unobserved factors are aggregated away. While concerns about unobserved factors at the county level still remain, panel data containing *repeated observations* at the county level enables the use of panel-data econometric techniques that can overcome such challenges. Furthermore, the monthly level of aggregation avoids the health outcome data being dominated by zeros, which is a problem with daily or weekly levels of aggregation. I aggregate each variable separately for the elderly (65+ years) and the younger adults (18-64 years). I then match the case surveillance data to the data on vaccination and add the data on population in 2019 by age group for each county from the U.S. Census Bureau. My balanced panel contains 199,950 observations (31 months * 2 age groups * 3,225 counties).

[8] Hortaçsu et al. (2021) document substantial under-reporting of infections in the early phases of the COVID-19 pandemic. While this concern applies only to infections, it is important to address. The use of rich fixed effects in TD and DiD designs substantially mitigates the risk that systematic under-reporting biases the estimated age-specific effects of vaccination on infections. These specifications absorb time-varying differences in reporting intensity across counties and ensure that identification relies on within-county variation across age groups and time.

4.2 Proxies for Health Outcomes

I employ the following proxies for the health outcomes related to COVID-19. First, I use total cases per million population to capture the incidence of COVID-19 in the county. I only count the confirmed cases in the CDC case surveillance data; however, all my results remain quantitatively similar when I include unconfirmed cases in my sample as well. Second, I examine the deaths per million population, which is my most significant proxy for health outcomes. I also similarly define ICU admissions and hospitalizations per million. Note that all these health outcome measures represent flow variables, i.e., the incidence of that particular health outcome in a county in a particular month.

4.3 Proxies for COVID-19 Vaccination

I use the *cumulative* percentage of the county population that has received two doses as my main proxy for the progress in vaccination against COVID-19. This metric is split into groups of the elderly and the younger adults. When examining the effect of the booster dose, I define a similar measure for the cumulative percentage of the county's population that has received the booster dose of the vaccine (again split into two groups by age).

4.4 Summary Statistics: Heterogeneous Impact of COVID-19 across Age Groups

Table 1 displays the summary statistics and Figure 3 displays the summary values of the health outcomes for each month from Jan-2020 to Jul-2022. As the data is organized at the (county, month) level, the averages and other summary statistics are computed at the (county, month) level. We observe the following empirical patterns:

1. Four waves of the pandemic are visible in the data with the deaths, ICU admissions, hospitalizations, and cases per million clearly showing four peaks in Apr-2020, Dec-2020, Aug-2021 and Jan-2022. The first peak for total cases per million is less easily visible, though it shows up in the data with the total cases per million briefly spiking in Apr-2020.

2. Deaths, especially among the elderly, were by far the highest in the first wave, at about 200 per million among the elderly and 50 per million among the younger adults. In each of the subsequent waves, deaths have been significantly lower, even though other health outcomes such as cases, ICU admissions, and hospitalizations were far greater.

3. The second wave, which occurred during the Omicron variant of the COVID-19 pandemic, was the most intense when the monthly averages for cases, ICU admissions, and hospitalizations were far greater than at any other time. However, deaths per million were lower in the second wave than in any of the other waves.

4. Deaths, ICU admissions, and hospitalizations per million are uniformly much higher for the elderly than for younger adults. The average deaths per million for the elderly is 4.3 times

higher than for younger adults (10.7 versus 2.5). The average ICU admissions per million for the elderly is 5.1 times higher than for younger adults (142 versus 28). Finally, the average hospitalizations per million for the elderly is 4.9 times higher than for younger adults (1328 versus 271). These disparities reflect the heightened risk of severe outcomes that the elderly face due to age-related vulnerabilities and comorbidities.

5. In contrast to these differences in deaths, ICU admissions and hospitalizations, total cases per million are almost identical for the elderly and younger adults. In Figure 3, we observe the two curves continually overlapping with each other for cases per million. Table 1 shows that the cases per million were 9% lower among the elderly as compared to the younger adults (2534 versus 2792). Infection rates for the two groups were almost identical during major national surges as well. The largest divergence—where younger adults' case rates significantly exceeded those of the elderly—occurred during the Omicron wave, consistent with the variant's higher transmissibility and changing behavioral patterns.

These descriptive patterns illustrate the *heterogeneous impact of COVID-19 across age groups.* The elderly consistently faced more severe health consequences—deaths, ICU admissions, and hospitalizations—while differences in infection rates were relatively smaller or higher among the younger adults. This underscores the potential for age-targeted vaccination strategies to yield differentiated returns: vaccinating the elderly may disproportionately reduce severe health outcomes, while vaccinating younger adults may have a larger effect on limiting disease transmission. These observed disparities motivate the empirical strategy used in the paper to quantify and compare vaccine effectiveness across age groups.

5 Empirical Strategy: Identifying Heterogeneous Returns from Vaccination

I exploit the age-based prioritization embedded in U.S. COVID-19 vaccination guidelines—along with substantial variation across counties and over time in vaccine uptake—to construct a quasi-experimental setting for identifying heterogeneous vaccination effects. As explained in section 3, elderly populations (65+) received vaccines earlier and more extensively than younger adults (18–65), creating natural differences in exposure that are sharp both cross-sectionally and temporally.

To rigorously quantify heterogeneity in vaccine returns across age groups, I adopt a *two-pronged empirical strategy.* First, I implement a *triple-difference* (TD) design that measures the relative (differential) impact of vaccination on the elderly versus the younger adults. Second, I separately estimate *difference-in-differences* (DiD) models for each age group to recover the absolute treatment effects within each group. Using these two approaches in concert allows for a full decomposition: the TD estimates provide causal evidence on which group benefits more per vaccine dose, while the DiD estimates reveal the magnitude of vaccine benefits for each group in levels.

5.1 Triple-difference design

My primary identification strategy is a TD design that leverages variation across three dimensions—age group, time, and vaccine roll-out intensity—to isolate the heterogeneous treatment effect of vaccination on outcomes. Intuitively, I compare how much outcomes improve for the elderly relative to the younger adults under differing degrees of vaccine uptake. Within this framework, *identification comes from differences in vaccination intensity between age groups in the same county and month.*

The TD estimator is estimated using a panel regression with a triple interaction term as the coefficient of interest: Elderly × Post-vaccine × High-intensity, or equivalently Elderly × VaccinationRate because VaccinationRate captures the effect of Post-vaccine × High-intensity. The estimating equation is:

$$y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn_{ic(t-12)} \cdot \Lambda_{65+,i} + \beta_2 Vaccn_{ic(t-12)} + \beta_3 \Lambda_{65+,i} + \varepsilon_{ict}, \tag{1}$$

where y_{ict} denotes a health outcome for age group i in county c in month t and $Vaccn_{ic(t-12)}$ denotes the cumulative proportion of the population of age group i that has been vaccinated in county c until *a year earlier*, i.e., by month $(t-12)$. $\Lambda_{65+,i}$ denotes a dummy that equals 1 for the elderly and 0 for the younger adults and β_c, and β_t denote county and time fixed effects respectively. The fixed effects for each county β_c soak up all the time-invariant variation due to observed and unobserved factors at the county-level, while the fixed effects for each sample month β_t enable us to control for inter-temporal variation in COVID-19 health outcomes that are common to all counties due to observed and unobserved factors. To alleviate concerns of auto-correlation, I cluster standard errors at the county-level.

The key coefficient of interest is β_1, which is effectively estimated as a triple-difference:

$$\textbf{TD} \equiv DiD|_{Elderly\ (65+)} - DiD|_{Younger\ Adults\ (18-64)}\,, \tag{2}$$

where the DiD is as defined as in equation (4), that is explained below. By differencing twice (across time and across counties) for each age group, then taking the difference between those differences estimated for each age group, the TD approach nets out any common trends and baseline differences, yielding a plausibly causal estimate of heterogeneity in vaccine returns. This is our most rigorous means of identifying heterogeneous effects, as it effectively uses the younger adults as a control group for the elderly (and vice versa) under the same conditions, akin to a "difference of a difference-in-differences." β_1 *is identified purely from differences in outcomes between age groups within the same county-month*, as vaccination rates change – significantly alleviating concerns of omitted variable bias at the county or time level.

The *identifying assumption* for the TD estimator is that in the absence of vaccination, the within-county gap in health outcomes between the elderly and the younger adults would have evolved similarly across different vaccine roll-out intensities. Even if a county accelerates vaccination in response to worsening health metrics, such policy responses would plausibly affect both the elderly

and younger adults within that county, and thus cancel out in the TD estimation. As long as the within-county allocation across age groups is not itself endogenously determined by unobservable shocks that affect one group but not the other, β_1 remains an unbiased estimate of the differential treatment effect. The identification thus hinges on a plausible quasi-experimental contrast in vaccine coverage between elderly and younger adults, conditional on shared local conditions in time.

5.2 Difference-in-difference estimates for elderly and younger adults

While the TD framework tells us how the effects differ between age groups, it does not directly estimate the absolute improvement in health outcomes due to vaccination for each group. I therefore also implement a standard difference-in-differences (DiD) design—as postulated in Imbens and Wooldridge (2009)—separately for each age group. For these DiD analyses, I compare outcomes before and after vaccination roll-outs between counties (or time periods) with different vaccination exposure, focusing on one age category at a time. The estimating equation I use is:

$$y_{ct} = \beta_c + \beta_t + \beta_1 \cdot Vaccn_{ic(t-12)} + \varepsilon_{ct}, \tag{3}$$

where β_c and β_t denote fixed effects for each county and sample month respectively and all other variables are as defined before. β_1 is estimated as a generalized DiD, as postulated in Imbens and Wooldridge (2009):

$$\textbf{DiD} \equiv (\bar{y}_{after} - \bar{y}_{before})|_{High\ Vaccination\ Counties} - (\bar{y}_{after} - \bar{y}_{before})|_{Low\ Vaccination\ Counties}, \tag{4}$$

where $\bar{y}_t$ denotes the mean estimated over time period t. Thus, β_1 estimates the difference in *within-county* before-after differences in vaccination in counties that reached higher levels of vaccination in a particular month vis-à-vis similar before-after differences in counties that did not reach similar levels of vaccination in the same month.

A redeeming feature of the county-level DiD estimation is that policy changes occur at the state level. So, the variation in vaccination at the county-level is likely to be quasi-random. In fact, Allcott et al. (2024) show that neither policy nor rates of voluntary social distancing explain a meaningful share of geographic variation of the COVID-19 pandemic; exogenous county-level characteristics predict such variation the best. As these characteristics are controlled for using county fixed effects, my DiD estimates are unlikely to be affected by county-level confounding factors.

5.3 Reverse Causality from Prior Outcomes?

Figure 4 examines whether early vaccination rates in Mar-2021 were systematically correlated with the severity of local COVID-19 outcomes in Dec-2020, i.e., whether reverse causality could plausibly explain the observed variation in vaccine uptake across counties. Note that even with potential reverse causality, the TD estimates are likely to be unbiased because of the comparison of the effects

between the younger adults and the elderly in the same (county, month). So, any reverse causality that affects both groups equally (e.g., county-wide policy responses to rising cases) is netted out by design. However, potential reverse causality can impact identification of the DiD estimates.

Each panel in Figure 4 plots the relationship between county-level COVID-19 health metrics in Dec-2020 on the x-axis and the share of each age group vaccinated by Mar-2021 on the y-axis, separately for individuals receiving at least one dose (blue circles) and two doses (red diamonds). The top two panels display the correlations for the elderly while the bottom two panels display the same for the younger adults. Across all panels, the visual evidence displays no systematic correlation between local COVID-19 severity and subsequent vaccine coverage. In fact, statistical tests show that none of the four correlations displayed across the four charts are statistically significant. In the top-left panel, counties with high elderly death rates in Dec-2020 are not consistently associated with higher elderly vaccination rates by Mar-2021. Vaccination uptake appears widespread and fairly uniform across a broad range of prior mortality experiences. Similarly, the top-right panel shows little evidence of a strong upward trend between case rates and vaccination for the elderly. While some high-incidence counties achieved moderate-to-high coverage, many others did not, and low-incidence counties often achieved similar uptake. The bottom panels present even flatter associations for younger adults. Most counties cluster at relatively low vaccination levels (under 40% with two doses), regardless of how severely affected they were in the preceding winter wave.

Together, these patterns suggest that vaccination coverage in Mar-2021 was not systematically driven by county-level disease burden in the preceding months, at least not in a way that would compromise the identification strategy. While some correlation might exist due to triaging of doses in severely affected regions, the scatter is broad and unstructured—supporting the claim that county-level variation in vaccination rates is not solely a function of endogenous response to prior outcomes. Thus, figure 4 alleviates concerns about reverse causality biasing the empirical analysis.

6 Results

6.1 Triple-difference estimates

In Table 2, I report the TD estimates using regression equations (1). Focusing on the coefficient of interest, β_1, shown in the shaded rows, we find that it is statistically significant at the 1% level of significance for each of the dependent variables. Furthermore, β_1 is negative for deaths, ICU admissions and hospitalizations per million but positive for cases per million. Thus, the results from the triple-differences confirm that even one year after two doses of the COVID-19 vaccine, deaths, ICU admissions and hospitalizations per million reduce *disproportionately more* among the elderly than the younger adults. However, the reduction in cases per million is greater among the younger adults than among the elderly as seen in the positive coefficient estimate of β_1.

Next, we note significant differences between the elderly and the younger population in the health outcomes as reflected in β_3, the coefficient of the dummy for elderly $\Lambda_{65+,i}$. The elderly are associated with greater deaths, ICU admissions, and hospitalizations per million than the younger

adults. However, cases per million is lower among the elderly than the younger adult population. Also, in contrast to β_1 and β_3, whose estimates are consistently significant across all the health outcomes, the estimate of β_2 is insignificant for ICU admissions and is opposite in sign to β_3 for hospitalizations and cases. Given this, the coefficients estimate using the DiD become critical to understand the effects separately for the two age groups.

6.2 Difference-in-difference estimates: Elderly versus younger adults

I now report in Table 3 the results of the difference-in-difference tests as captured in regression equation (3). Panel A reports the results for the elderly while Panel B reports the same for the younger adults. Across both panels, I find that all of the coefficient estimates are negative and are significant at the 1% level of significance. When I compare the coefficient estimates between Panels A and B, I find that the impact one year after two doses of the vaccine on the elderly vis-à-vis the younger adults is 3.4 times (=0.24/0.07) for deaths per million, 13 times for ICU admissions per million, and 8 times for hospitalizations per million. However, the impact for cases per million among the elderly is one-third that of the impact on the younger adults. Overall, the effects of vaccination one year later on deaths, ICU admissions and hospitalizations being larger for the elderly than the younger adults and the same effect on infections (cases) being lower for the elderly than the younger adults are consistent with those observed in Table 2.

6.3 Dynamic effects and robustness of difference-in-differences

My difference-in-difference specification corresponds to a continuous treatment that is staggered in time across counties. Recent research shows that the identifying assumptions required in such setups are stronger than the canonical two group, two period setup. In particular, an additional identifying assumption is that the "average treatment effect function" does not vary with the percentage of the county population that is vaccinated (Sun and Abraham (2021) , De Chaisemartin and d'Haultfoeuille (2020), Callaway and Sant'Anna (2021), Roth et al. (2022)). To validate my difference-in-difference design, I use the methodology developed by Rios-Avila et al. (2022) to implement the robust estimation of difference-in-difference in our setup. Specifically, I use counties that have not yet been treated to serve as the control group; I regard the month in which vaccination begins in a county as the month of treatment. As the timing of the vaccination was different for the elderly and younger adults, I show the dynamic effects separately for both these groups, thereby allowing for the differences in the timing of treatment across the two groups. I show the effects from five months before treatment until five months after treatment. Figure 5 shows these dynamic effects, where 0 corresponds to the month of treatment when the percentage of population that has been administered two doses of the vaccine becomes non-zero in the county for the first time.

We note four findings that emerge from these plots. First, focusing on the dynamic effects for deaths, we notice that while the ATT estimates are strong for the elderly, the effects are relatively muted for the younger adults, which is consistent with the results seen so far and with the higher immunity possessed by the younger adults. Next, focusing on hospitalizations and ICU admissions,

we find that the ATT estimates are an order of magnitude higher for the elderly when compared to the younger adults. This aspect is also consistent with the results we have obtained so far and with elderly being more vulnerable to adverse health outcomes than the younger adults.

Second, the pre-treatment ATT estimates for all outcomes show a monotonic upward trend in the periods leading up to treatment: from –3 to –1 for cases and hospitalizations, –4 to –1 for ICU admissions, and from –3 to 0 for deaths. These patterns indicate that treated counties were on a worsening trajectory relative to the control group prior to vaccination. However, what matters for causal identification is not the level or direction of these differences, but *whether the relative trend would have remained stable in the absence of treatment.* In this case, the *sharp inflection and sustained downward* shift in all outcomes beginning at or immediately after treatment suggests a *discontinuity in trend, consistent with a treatment effect rather than a continuation of prior dynamics.* If vaccination had not occurred, we would expect the ATT to either stabilize or rise further; instead, it reverses direction in a way that coincides with the onset of vaccination. This pattern supports the parallel trends assumption as a plausible counterfactual trajectory—particularly when coupled with the fact that the reversal is consistent across outcomes and the age groups.

Third, the timing of treatment effects varies across outcomes in a way that reflects biological and clinical mechanisms. For cases, hospitalizations, and ICU admissions, the effects begin manifesting in the first month after vaccination, coinciding with the time required for vaccines to generate immune protection and reduce community spread. Notably, effects on deaths emerge only from the second month after vaccination, reflecting both the longer time from infection to mortality and potential delays in death reporting. This *staggered onset of benefits* strengthens the interpretation that the observed effects reflect the natural epidemiological unfolding of vaccine impact across the severity spectrum.

Finally, there are substantial and sustained post-treatment reductions in all outcomes for the elderly, with effects growing steadily over time. We notice that the decline in cases, hospitalizations and ICU admissions are significant from the first month after vaccination. In contrast, the decline in deaths is only significant from the second month after vaccination. Thus, cases, hospitalizations, and ICU admissions per million begin falling in the first month post vaccination while deaths per million begin falling only in the second month post vaccination. These dynamics closely match the known clinical efficacy of vaccines, particularly the protection it offers for high-risk individuals against severe/ fatal health outcomes.

Collectively, the ATT trajectories exhibit *coherent directionality, biological plausibility, and internal consistency* across both age groups and outcomes. Taken together, these features provide strong support for the claim that vaccination generated meaningful and sustained health improvements—particularly for the elderly—consistent with the core findings of the TD and DiD analyses.

6.4 Importance of examining effects one-year after

To further validate the dynamic treatment effects and assess the appropriate temporal window over which vaccine returns manifest, Table 4 presents regression estimates using varying lags between

health outcomes and vaccination—ranging from the same month (contemporaneous effects) to one year prior. Panel A reports triple-difference estimates, while Panels B and C show separate difference-in-differences results for the elderly and younger adults, respectively.

The results provide three important insights. First, *the sharp reversal in sign* and magnitude of estimates between shorter lags and the one-year lag validates the main results. For example, in Panel A, the coefficients for deaths and hospitalizations are statistically insignificant or even positive when using three-month or six-month lags, but become large and significantly negative with a one-year lag. This temporal reversal mirrors the findings we noted from the dynamic treatment effect plots and reinforces the plausibility of a causal interpretation.

Second, the DiD estimates offer further reassurance. In Panel B, for the elderly, the sign and size of effects also grow over time, with one-year-lag estimates consistent with those from the triple-difference models. For younger adults (Panel C), the one-year-lag coefficients are either smaller or reversed in sign—particularly for infections—highlighting that the benefits of vaccination differ systematically across age groups and outcomes, consistent with the public-good versus private-return framing introduced earlier.

Finally, the dynamic treatment effects reveal that the timing of vaccine benefits differs markedly across outcomes: reductions in infections and hospitalizations appear within the first month while the effects on mortality materialize only by the second month. Here, we must note that there is no inconsistency between the positive coefficients observed for shorter lags in Table 4 and the month-by-month dynamic revealing in Figure 5 that the effects start manifesting within a month or two after vaccination. The effects studied for different lags in Table 4 *aggregate* the effects shown month-by-month in Figure 5 and therefore reveal the comprehensive effect that is not captured in Figure 5.

These staggered dynamics underscore the importance of examining longer-run impacts of vaccination, beyond the short-term windows typically studied in the existing literature. By focusing on vaccine coverage twelve months earlier, the main analysis captures the cumulative and durable effects of vaccination—after immunity has had time to build, behaviors have adjusted, and downstream outcomes such as death have fully manifested. This long-horizon perspective provides a more complete view of vaccine effectiveness, particularly for severe outcomes that lag infection, and is especially relevant for informing resource allocation during future pandemics. In this respect, the *use of a one-year lag is not merely a robustness feature but a central conceptual aspect of this study.*

Overall, this analysis confirms that the long-run effects of vaccination are both larger in magnitude and more policy-relevant than short-run estimates, further justifying the paper's focus on one-year-lagged vaccine exposure. The temporal consistency across identification strategies strengthens confidence in the robustness and causal validity of the results.

7 Robustness Tests And Alternative Explanations

7.1 Effects of two doses versus the booster dose

As we are studying the effects one year after two doses of the COVID-19 vaccine are administered, a key concern that arises is the following. Are the effects of the two doses of vaccine being confounded by the effect of the booster dose administered later? In other words, could the effects of the vaccine given one year before actually reflect the effects of the booster dose? *This question cannot be answered using clinical trials data* because answering this question requires variation among individuals along two different margins: (i) among individuals who received two doses of the vaccine, those that did and did not receive the booster dose; and (ii) among individuals who received the booster dose of the vaccine, those that did and did not receive two doses of the vaccine. The first sample would help answer the incremental effect of the booster dose over and above the effect of the two doses, while the second sample would help answer the incremental effect of two doses of the vaccine when compared to the effect of the booster dose. The second sample, however, cannot be constructed by definition as only those individuals that have received two doses of the vaccine are eligible to get the booster dose. Therefore, this question can only be studied using the observational data as in this study.

Table 5 shows results of the tests of the following panel regression:

$$\begin{aligned} y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn2Doses_{ic(t-12)} \cdot \Lambda_{65+,i} + \beta_2 VaccnBooster_{ict} \cdot \Lambda_{65+,i} \\ + \beta_3 Vaccn2Doses_{ic(t-12)} + \beta_4 VaccnBooster_{ict} + \beta_5 \cdot \Lambda_{65+,i} + \varepsilon_{ct}, \end{aligned}$$

where I add the interaction of the dummy variable $\Lambda_{65+,i}$ with $VaccnBooster_{ict}$, which denotes the cumulative proportion of the population that has received the booster dose in county c by month t. Thus, the coefficient β_1 captures the incremental effect of two doses of the vaccine when compared to the effect of the booster as a TD estimate. I notice in Table 5 that all the coefficient estimates for β_1 remain negative and statistically significant. In fact, the coefficient estimates for β_2 are opposite those of β_1. As the administration of the booster dose may not have fully stabilized in several counties, the inference of the effects of the booster dose are quite tenuous. Therefore, for our purposes, the point of focus in these results is not the effect of the booster dose but showing that the effect of the two doses of the vaccine are not due to the effect of the booster dose. Thus, Table 5 provides robust evidence that the effect of the two doses of the vaccine cannot be attributed to the effect of the booster dose.

7.2 Accounting for spatial spillovers: MSA level analysis

A potential concern with analysis at the county level is that *some people may get vaccinated and/or hospitalized in a neighboring county.* To address concerns that my results are affected by such concerns, I aggregate my data at the level of MSA and re-run the tests of regressions **??** and 3. This aggregation helps to internalize most cross-county mobility within metro areas and includes

the effects of most proximate spillovers. Table 6 shows the results of this analysis. I find that the results for the triple-difference tests shown in Panel A and the difference-in-difference tests for elders and younger adults in Panels B and C respectively remain qualitatively similar. While the coefficients for the difference-in-difference estimates for hospitalizations and cases among elders are not statistically significant, their signs remain the same. Similarly, the coefficients for the difference-in-difference estimates for ICU admissions and hospitalizations among younger adults are not statistically significant. However, as the coefficients in the triple-difference tests in Panel A are all uniformly significant at the 1% level, it is clear that the disproportionate effects of two doses of the vaccine even one year later remains unchanged when we analyze the effects at the MSA level instead of the country level. In fact, by comparing with the estimates obtained in Table 2, we notice that when we analyze by aggregating at the MSA level, the magnitude of the coefficients is substantially larger for deaths, ICU admissions and hospitalizations while the coefficient for infections is slightly lower.

7.3 Controlling for other confounders using different sets of fixed effects

Are the main results presented in Table 2 driven by the specific fixed effects structure employed in the baseline specification? This specification absorbs two broad classes of omitted variables: (i) county-level time-invariant factors such as demographic composition, health infrastructure, political preferences, and baseline exposure risks; and (ii) nationwide or calendar-time shocks such as the arrival of a new variant, national vaccination campaigns, or federal policy shifts, which affect all counties simultaneously. To address other classes of confounders, we estimate the TD with alternative sets of fixed effects.

7.3.1 Triple-difference with (county, time) fixed effects

The baseline TD design compares the DiD estimate for the elderly to the corresponding estimate for the young within the same county and month. This comparison already nets out shocks that are common to both groups in that county-month. With county and month fixed effects included, the specification also absorbs time-invariant county characteristics (such as demographics, baseline health infrastructure, or persistent differences in health behaviors) and national-level shocks that vary across months (such as the emergence of new variants or changes in federal guidance).

Nevertheless, there remain county-specific, time-varying confounders that can bias the TD estimates despite these controls. For example, consider a county that experiences a localized outbreak in a nursing home. Such an outbreak raises infections among the elderly much more than among the young in that month, even if vaccination coverage has not changed. A simple TD with county and month fixed effects would attribute part of that elderly-specific spike to differential vaccine effectiveness, rather than to the outbreak itself. Similarly, suppose a county government implements a policy that targets one age group more than the other—such as closing schools (affecting mostly younger adults) or restricting visits to eldercare facilities (affecting mostly the elderly). These county-level policies change exposure risk in an age-differentiated way, potentially biasing

the TD estimates if not fully controlled. Other examples include age-specific shocks to health-seeking behavior (e.g., a county hospital running a campaign for vaccine uptake among the elderly) or differences in testing intensity across age groups that vary by county and month.

By including fixed effects for each (county, time) pair, we absorb all county-specific shocks in each month, whether they affect age groups equally or unequally. This stringent specification ensures that identification comes solely from within-county-month variation across age groups, effectively comparing the elderly and younger adults under shared local conditions. The key regressor—the interaction of vaccination rate and the elderly dummy—varies at the (county, time, age group) level, so its coefficient is identified from differences in outcomes between age groups within the same county and month, after all local shocks have been netted out. This approach eliminates confounding from county-month-specific factors such as: (i) local epidemic dynamics (e.g., super-spreading events); (ii) county-wide policies (e.g., mask mandates or business closures); (iii) environmental or behavioral shifts (e.g., weather, mobility, or hospital strain); and (iv) changes in testing or reporting practices.

Panel A of Table 7 presents results from this specification. We notice that compared to the baseline specification in Table 2, the TD coefficient estimate for deaths, ICU admissions and hospitalizations is at least as much in Panel A of Table 7. However, the TD coefficient estimate for cases is lower in Panel A of Table 7.

The positive coefficient estimate for the average effect of vaccination here may seem counter-intuitive. We note that during periods like the Omicron wave, cases, hospitalizations, and even deaths surged across the board while vaccination rates were still increasing. The (county, time) FE possibly manifests this feature in a way that the county and time FE in Table 2 may not capture. This is precisely why the differential effect across age groups (as estimated via the TD) is the key parameter of interest.

7.3.2 Triple-difference with (county, age group) and time fixed effects

A residual concern, even with the inclusion of county-by-month fixed effects, is that while it accounts for all time-varying shocks at the county level such as local COVID waves, lockdowns, or hospital strain, the identifying variation could still be *confounded by age-specific behavioral or policy responses that differ within counties over time*. For example, during the pandemic, elderly individuals may have responded differently than younger adults—being more cautious in social interactions, adhering more strictly to mask mandates, or altering health-seeking behavior in ways that correlate with vaccination rates. If these behavioral differences evolved over time and were correlated with vaccination uptake, they could bias even the TD estimated with county-time fixed effects.

To address this concern, I re-estimate the TD using alternative fixed effects specifications that more directly account for age-group-specific shocks or behaviors. Panel B of Table 7 presents results from a specification that includes fixed effects for each (county, age group) pair, along with time fixed effects. This specification absorbs any *time-invariant differences between elderly and younger adults*

within a county. Such differences may include baseline health conditions (e.g., higher comorbidities among the elderly), age-specific healthcare access and utilization (e.g., Medicare coverage, hospital admission thresholds), systematic differences in testing behavior, or even long-standing social and occupational exposure patterns. For instance, younger adults are more likely to be employed in high-contact service jobs, while older adults may seek healthcare more readily when symptomatic. By assigning a separate intercept to each county–age group pair, the specification nets out all of these persistent age-group disparities within counties. Time fixed effects then absorb all shocks that affect both counties and age groups equally at a national scale, such as the emergence of new variants or federal-level policy changes. The results from this specification remain quantitatively similar to the baseline estimates in Table 2, suggesting that the core findings are not driven by time-invariant differences in age-specific county characteristics.

7.3.3 Triple-difference with state-specific time trends as well as county and time fixed effects

Next, we estimate models with county and time fixed effects together with state-specific linear time trends. Panel C of Table 7 reports results from this specification, which absorbs *gradual, state-level changes in epidemic dynamics and policy environments that unfold differently across states*. These may include the state-specific public health mandates such as masking and reopening schedules, differences in the speed of variant penetration across states, or smooth shifts in mobility and behavior linked to state-level political or cultural factors. For example, some states reopened schools and businesses gradually over several months, while others maintained restrictions longer, creating smooth but heterogeneous trajectories in transmission risk. By allowing each state to have its own time trend, the specification nets out these slow-moving, state-specific policy and epidemic dynamics that might otherwise be correlated with vaccination uptake. This ensures that the estimated threshold effects are not simply picking up broader state-level trends in cases or deaths. We see that the results remain strong and if anything become stronger.

In all these complementary TD designs, the results remain strong and closely aligned with our baseline TD estimates. This robustness is crucial because each specification rules out a distinct class of omitted variables: the (county, age group) fixed effects absorb persistent age-specific differences within counties, while the state-specific time trends absorb gradually evolving, state-specific shocks. Together with the TD specification that uses (county, time) fixed effects to absorb shocks common to both age groups within a county-month, the convergence of results across these demanding specifications provides compelling evidence that our findings are not driven by unobserved confounders. The consistency across designs strengthens the credibility of the causal interpretation: heterogeneous returns to vaccination reflect real differences in the private and public benefits of immunization, not artifacts of omitted variables.

Taken together, these robustness checks lend support to the main findings. These alternative fixed effects structures demonstrate that the core estimates are not driven by specific modeling choices.

7.4 Analysis after dropping observations in 2020

A concern could be that my results are driven possibly by the zero values for vaccination in the year 2020 when COVID-19 vaccination had not begun. To address these concerns, I drop the observations for the year 2020 and redo the analysis shown in Tables 2 and 3. Table 8 displays the results for this robustness analysis. The results for the triple-difference tests are shown in Panel A and the difference-in-difference tests for elders and younger adults are shown in Panels B and C respectively. We note that the results remain qualitatively and quantitatively similar to those observed in Tables 2 and 3. Therefore, we can infer that the results are unlikely to be driven by the zero values for vaccination in the year 2020.

8 Estimating Economic Effects

8.1 Discussion on stability of coefficients

To evaluate the robustness and magnitude of estimated vaccine effects across model variations, Panel A of Table 9 summarizes the TD coefficients on deaths, ICU admissions, hospitalizations, and confirmed cases under alternative specifications. These include models with different fixed effects (e.g., county and time fixed effects separately, county-time fixed effects, county-by-age and time fixed effects, state-specific time trends), estimation at the MSA level, controls for the booster dose, and a specification that excludes the observations in the year 2000. Across all outcomes, the direction and statistical significance of the effects remain consistent, confirming the strength of the identification strategy. Notably, while the magnitudes of the estimates vary—often becoming larger in absolute terms when more granular controls or MSA-level aggregations are used—the original estimates from Table 2 consistently lie at the lower bound of the plausible range. This suggests that the original specification with only county and only time fixed effects are the most conservative across the suite of models considered.

Having inferred that the baseline TD model provides estimates that are robust, we can use the same to compare the same with the results obtained from the DiD design. In Panels B and C, we compare the coefficients from the DiD model estimated in Table 3 with the coefficients derived from the baseline TD model in Table 2. We note that the DiD estimates are not only along the expected lines, they are very close to the TD estimates in most cases.

Recall that in section 5.2, we had discussed that a redeeming feature of the county-level DiD estimation is that policy changes occur at the state level. So, the variation in vaccination at the county-level is likely to be quasi-random. Moreover, Allcott et al. (2024) show that neither policy nor rates of voluntary social distancing explain a meaningful share of geographic variation of the COVID-19 pandemic; exogenous county-level characteristics predict such variation the best. As these characteristics are controlled for using county fixed effects, the DiD estimates are unlikely to be affected by county-level confounding factors. Taken together, by summarizing the coefficient estimates across several empirical models, we are able to reinforce the credibility and stability of the

core empirical findings while indicating that the baseline TD and the DiD estimates are conservative in magnitude. Therefore, we can infer that the DiD estimates are the most appropriate to use for estimating the economic magnitudes.

8.2 Economic magnitudes

I now assess the economic magnitudes of the heterogeneous effects. I use the averages for the various health outcomes reported in Table 1 separately for the elderly and for younger adults along with the effects of the two doses of the vaccine one year before as estimated in Table 3 to assess the economic magnitudes of the effects.

A 10% increase in the percentage of the elderly population vaccinated in a county leads into a decrease of 2.4 deaths per million, 22.4 ICU admissions per million, 80.4 hospitalizations per million, and 80 cases per million. In contrast, a 10% increase in the percentage of the younger adult population vaccinated in a county leads into a decrease of 0.7 deaths per million, 1.7 ICU admissions per million, 10.3 hospitalizations per million, and 240 cases per million. Thus, compared to the younger adults, vaccinating 10% more of the elderly population in a county reduces deaths 3-fold more, ICU admissions 13-fold more and hospitalizations 8-fold more. Conversely, compared to the elderly, vaccinating 10% more of the younger adult population in a county reduces infections 3-fold more.

These absolute magnitudes can be complemented by estimates of relative magnitudes as well. For the elderly, a one standard deviation increase in the percentage of the elderly population vaccinated in a county (= 37.5 as seen in Panel A of Table 1) leads into a decrease of 9 deaths per million, 84 ICU admissions per million, 301 hospitalizations per million, and 300 cases per million. I compare these effects to the averages for the elderly as seen in Panel A of Table 1 of 10.7 deaths per million, 683 ICU admissions per million, 1328 hospitalizations per million, and 2534 cases per million. Thus a one standard deviation increase in the percentage of the elderly population vaccinated in a county causes a decrease of 84% in deaths per million, 12% in ICU admissions per million, 23% in hospitalizations per million, and 12% in cases per million. Clearly, the economic effect of vaccination even one year later remains large and is most pronounced for the most debilitating outcome: deaths. The economic effect also remains much lower but meaningful for the other health outcomes.

For younger adults, a one standard deviation increase in the percentage of the population of younger adults vaccinated in a county (= 27.1 as seen in Panel B of Table 1) leads into a decrease of 1.9 deaths per million, 4.6 ICU admissions per million, 27 hospitalizations per million, and 653 cases per million. I compare these effects to the averages for the younger adults as seen in Panel B of Table 1 of 2.5 deaths per million, 28 ICU admissions per million, 271 hospitalizations per million, and 2792 cases per million. Thus a one standard deviation increase in the percentage of the population of younger adults vaccinated in a county causes a decrease of 75% in deaths per million, 16% in ICU admissions per million, 10% in hospitalizations per million, and 23% in cases per million. Even for the younger adults, the economic effect of vaccination even one year after

remains large and is most pronounced for the most debilitating outcome and remains much lower but meaningful for the other health outcomes.

9 Counter-factual Estimates

To evaluate the policy implications of age-targeted vaccination strategies, I simulate three distinct counterfactual scenarios based on alternative allocations of COVID-19 vaccines across age groups:

1. Directing all vaccines exclusively to the elderly until they are fully vaccinated,
2. Directing all vaccines exclusively to younger adults until they are fully vaccinated, and
3. Allocating vaccines proportionally to each group's share in the county's population.

In each case, I hold the total monthly vaccine supply in each county fixed at its actual level and re-allocate doses across age groups accordingly. Using the estimated treatment effects as summarized in Panel B of Table 9, I compute predicted monthly *totals* (not per-capita rates) for deaths, ICU admissions, hospitalizations, and infections across both age groups one year later. These counter-factual estimates quantify how different allocation rules would have affected pandemic outcomes under realistic logistical constraints. As we have data till

9.1 Vaccinating Only the Elderly

Panel A of Figure 6 displays the simulated outcomes if all vaccine doses were allocated exclusively to the elderly population until full coverage was reached. Because younger adults became broadly eligible only by Mar-21, this counterfactual begins in Mar-22 to enable a consistent comparison. To aid comparison, I report the average across all counties of the health outcomes per million of population. All the outcomes are The results show large, compounding gains in severe outcomes among the elderly. Compared to the real outcomes, on average per county per month, deaths decline additionally by about 8 per million, ICU admissions by nearly 75 per million, and hospitalizations by about 250 per million by Jul-22. By comparing with the sample averages reported in Table 1, we see that these counter-factual outcomes represent 75% (=8/10.7), 53% (=75/142), 19% (=250/1327) of the average deaths, ICU admissions and hospitalizations per million per county per month. The 75% decline in deaths and 53% decline in ICU admissions represents a significant improvement. While the 19% decline in hospitalizations is not trivial either, the decline is not as pronounced as the outcomes for deaths and ICU admissions. Consistent with higher immunity among the younger adults, the increases in these health outcomes for the younger adults are small. However, reflecting the cost of neglecting transmission control, infections rise substantially among younger adults by nearly 270 per million on average per county, which reflects around 10% (=270/2792) increase in cases compared to the average per million per county. This is offset by the decline in cases of about 250 per million among the elderly, which reflects a 4.6% (=250/2534) decrease in cases per million per county.

9.2 Vaccinating Only Younger Adults

Panel B of Figure 6 presents the opposite scenario, where all vaccines are directed to younger adults starting in Jan-22. To aid interpretation, I report the average across all counties of the incremental health outcomes per million population.

The results show large, compounding gains in health outcomes among younger adults, but significant adverse effects among the elderly. Compared to the real outcomes, on average per county per month, deaths increase by about 14 per million among the elderly by Jul-22, while declining slightly among younger adults. ICU admissions rise by approximately 130 per million among the elderly and remain nearly flat for younger adults. Hospitalizations increase by roughly 470 per million among the elderly, while decreasing slightly among younger adults. Infections decline substantially among younger adults—by nearly 400 per million on average per county—reflecting decreased transmission. This is offset by an increase of about 420 per million in infections among the elderly.

By comparing these results with the sample averages reported in Table 1, we see that the increase in deaths among the elderly represents about 130% (=14/10.7) of the average monthly death rate per million, and the rise in ICU admissions amounts to 92% (=130/142) of the average rate. The surge in hospitalizations among the elderly is about 35% (=470/1327). Meanwhile, the rise in infections among younger adults accounts for about 14% (=400/2792) of the average case rate, while the reduction in cases among the elderly reflects a 16% (=420/2534) decrease from the actual outcome.

These findings highlight the substantial trade-offs involved in prioritizing vaccination for younger adults: while it reduces infections and mild disease in this group, it leads to significantly higher morbidity and mortality among the elderly, who remain vulnerable without protection.

9.3 Proportional Allocation Based on Population Shares

Panel C of Figure 6 reports results from a proportional allocation strategy, in which each county's actual vaccine supply is distributed across age groups in line with their population shares. The outcomes from this policy fall between the two extremes. Compared to the actual rollout, this policy leads to moderate increases in elderly health burdens: deaths rise by about 6 per million (56% of average monthly deaths), ICU admissions by 55 per million (39%), and hospitalizations by 200 per million (15%). Infections fall by approximately 150 per million among younger adults (5.4% of average), reflecting improved transmission control. This represents a middle ground—better than prioritizing younger adults, but worse than focusing on the elderly—highlighting the trade-off between protecting the vulnerable and maintaining broader population immunity.

9.4 Policy Implications

Taken together, these simulations reveal sharp trade-offs in the effects of age-based vaccine allocation. The younger-first strategy maximizes infection reduction but generates substantial avoidable

mortality and hospitalization among the elderly. The elderly-first strategy yields the strongest protection against severe outcomes, while infections rise moderately among the unvaccinated younger group. The proportional allocation policy offers a compromise: it achieves moderate reductions in both mortality and infections, avoiding the worst-case outcomes of either extreme.

While the counterfactual outcomes may seem to suggest that focusing the vaccination on the elderly may be the optimal policy response, we must note that the economic impact of a 10% rise in infections among the younger adults may be large and could potentially even offset the gains for the elderly. For example, a 10% rise in infections would translate into similar absenteeism at work, leading to negative impacts on production, supply chains, and thereby on aggregate supply. The reduced supply would, in turn, impact inflation and output, which would, in turn, lead to another round of economic effects on unemployment and economic distress. Similarly, the economic benefits of 14% lower infection among the younger adults when the policy option to vaccinate the young first adds to the benefits of this policy option and must be factored in when assessing this scenario. However, if we remain focused only on the health outcomes, the results suggest that the *elderly-first strategy dominates*—particularly if the social planner places high weight on reducing deaths and ICU burdens. While the proportional strategy offers political or ethical appeal because of the equity it achieves across groups, it under-performs in terms of life-saving outcomes. These findings support sequential or hybrid allocation frameworks in future pandemic settings: prioritize vulnerable groups first to mitigate severe health outcomes, and then shift focus to transmission control among lower-risk but socially active populations.

10 Discussion: Broader Lessons for Pandemic Policy

The findings of this study speak directly to a central dilemma in pandemic policy: how to allocate scarce vaccines when their private and public returns diverge across population subgroups. While the empirical analysis leverages rich U.S. county-level data, the underlying trade-off between protecting the most vulnerable and minimizing disease transmission is conceptually general and likely to arise in most epidemiological settings. Nevertheless, extrapolating the precise quantitative magnitudes to other countries requires careful consideration of demographic, institutional, and healthcare-specific contexts.

10.1 Application to other country settings

The U.S. population is older than many low- and middle-income countries but younger than several European or East Asian societies. In countries with younger populations, the social returns to vaccinating the working-age cohort might be relatively higher, amplifying the public-good rationale for targeting those groups. Conversely, countries with more robust healthcare systems may be better able to mitigate the burden on ICU and hospitalization capacity, potentially dampening the marginal benefit of vaccinating high-risk elderly groups. Similarly, the nature of community transmission—driven by household size, urban density, mobility patterns, and occupational expo-

sure—can affect the externalities from vaccinating younger adults. Taken together, these patterns suggest that the core insight of this study—heterogeneous returns to vaccination differ based on epidemiological risk—holds broadly across settings, with optimal prioritization requiring modest adjustments to reflect local demography, health system capacity, and transmission dynamics.

10.2 Applicability beyond COVID-19

The period covered in this study spans several COVID-19 variants, including Alpha, Delta, and the Omicron wave. Though vaccine efficacy varied across variants, the empirical strategies employed—especially the triple-difference design—are designed to isolate average treatment effects under the prevailing epidemiological environment. Nonetheless, future pandemics may differ in their age-specific severity or transmission dynamics. For instance, a pathogen that disproportionately affects younger adults would invert the policy calculus examined here. What carries across settings is the logic of quantifying both private protection and public-good spillovers; the specific targeting rule can then be re-calibrated in real time as new pathogens emerge and scientific evidence accumulates.

10.3 Macroeconomic dimensions of the vaccination trade-off

From an economic standpoint, the divergent effects observed in this study have meaningful implications for macroeconomic policy. The sharp reduction in infections among younger adults—those most likely to participate in the labor force—likely translated into measurable gains in labor supply, reduced absenteeism, and higher productivity. While this paper does not estimate those channels directly, the connection between population health and economic output is well-established (Perkins et al., 2023). Recent research shows that even modest reductions in transmission can generate sizable economic dividends, especially when targeted toward economically active populations (Deb et al., 2021). For instance, fewer infections in working-age cohorts could sustain consumer-facing industries, reduce strain on social insurance programs, and lower medical expenditure. The paper's findings on reduced severe health outcomes among the elderly also suggest potential fiscal benefits by avoiding costly ICU and long-term care utilization. Together, these mechanisms highlight that the vaccination trade-off is not merely clinical; it is also profoundly economic.

In sum, this study underscores the importance of designing vaccination policy with both *heterogeneous biological returns* and *public externalities* in mind. The findings offer not only a template for interpreting COVID-19 vaccination outcomes but also a generalizable framework for future pandemic preparedness—one that integrates epidemiological risk, demographic structure, healthcare capacity, and economic spillovers into the calculus for designing the optimal vaccination policy.

11 Conclusion

This paper provides new empirical evidence on the heterogeneous benefits of COVID-19 vaccination across age groups, highlighting a central trade-off in pandemic policy design: whether to

prioritize protecting the most clinically vulnerable or curbing transmission through socially active populations. Using population-level panel data from all U.S. counties and two complementary quasi-experimental designs—a triple-difference framework and difference-in-differences estimation—I show that vaccines generate sharply different returns by age group and health outcome.

The results demonstrate that vaccinating the elderly yields far greater reductions in deaths, ICU admissions, and hospitalizations, while vaccinating younger adults more strongly reduces infections, reflecting their central role in transmission. These findings are robust across alternative specifications and outcomes. Together, they reveal that marginal returns to vaccination differ not only in magnitude but in kind: doses given to the elderly save lives directly, while doses given to the young primarily curb spread.

Counterfactual simulations inform future pandemic policy. An elderly-first strategy minimizes mortality and pressure on critical care systems; a younger-first strategy curtails infections but leads to substantial avoidable deaths; and proportional allocation delivers intermediate results without dominating on either margin.

Beyond their immediate policy relevance, the findings speak to foundational questions in the economics of health and crisis management. They demonstrate how heterogeneity in clinical vulnerability and social behavior mediates the returns to preventive interventions, and they highlight that vaccines are not only biomedical tools but also economic instruments whose effectiveness depends critically on allocation. Methodologically, the paper shows how quasi-experimental variation in population-level data can generate real-time, policy-relevant insights in emergency contexts.

Looking forward, several extensions are promising. Embedding these causal estimates into structural models could yield formal characterizations of optimal allocation over time and space. Linking vaccine-induced reductions in infections among working-age adults to labor supply and productivity would connect micro-level health effects to macroeconomic performance. Extending this framework to other countries with different demographic structures and health system capacities would provide valuable evidence on how the inter-generational trade-off varies across institutional settings, informing global vaccination strategies.

Taken together, the study demonstrates how population-level data can be used to measure both the private and public returns to vaccination, and how these estimates can guide the design of vaccination campaigns and emergency response strategies in future pandemics.

References

Aassve, A., T. Capezzone, N. Cavalli, P. Conzo, and C. Peng (2024). Social and political trust diverge during a crisis. *Scientific Reports 14*(1), 331.

Acemoglu, D., V. Chernozhukov, I. Werning, and M. D. Whinston (2021). Optimal targeted lockdowns in a multi-group SIR model. *American Economic Review: Insights 3*(4), 487–502.

Adda, J. (2016). Economic activity and the spread of viral diseases: Evidence from high frequency data. *The Quarterly Journal of Economics 131*(2), 891–941.

Allcott, H., L. Boxell, J. Conway, B. Ferguson, M. Gentzkow, and B. Goldman (2024). What explains temporal and geographic variation in the early us covid-19 pandemic? *Review of Economic Design*, 1–41.

Almond, D., J. J. Doyle Jr, A. E. Kowalski, and H. Williams (2010). Estimating marginal returns to medical care: Evidence from at-risk newborns. *The Quarterly Journal of Economics 125*(2), 591–634.

Alvarez, F. E., D. Argente, and F. Lippi (2021). A simple planning problem for covid-19 lockdown, testing, and tracing. *American Economic Review: Insights 3*(3), 367–82.

Andersson, O., P. Campos-Mercade, A. N. Meier, and E. Wengström (2021). Anticipation of COVID-19 vaccines reduces willingness to socially distance. *Journal of Health Economics 80*, 102530.

Brown, J. R. and A. Finkelstein (2008). The interaction of public and private insurance: Medicaid and the long-term care insurance market. *American Economic Review 98*(3), 1083–1102.

Bundorf, M. K., J. DeMatteis, G. Miller, M. Polyakova, J. L. Streeter, and J. Wivagg (2022). Risk perceptions and private protective behaviors: Evidence from COVID-19 pandemic. *The Review of Economics and Statistics*, 1–40.

Cakmakli, C., S. Demiralp, S. Kalemli-Ozcan, S. Yesiltas, and M. A. Yildirim (2021). The economic case for global vaccinations: An epidemiological model with international production networks. Technical report, National Bureau of Economic Research.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Card, D., C. Dobkin, and N. Maestas (2008). The impact of nearly universal insurance coverage on health care utilization: evidence from medicare. *American Economic Review 98*(5), 2242–2258.

Chang, R., H. Martínez, and A. Velasco (2021, April). Pandemics, incentives, and economic policy: A dynamic model. Working Paper 28636, National Bureau of Economic Research.

Chetty, R., J. N. Friedman, and M. Stepner (2024). The economic impacts of covid-19: Evidence from a new public database built using private sector data. *The Quarterly Journal of Economics 139*(2), 829–889.

Cutler, D. M. and J. Gruber (1996). Does public insurance crowd out private insurance? *The Quarterly Journal of Economics 111*(2), 391–430.

Davies, W. and X. Zhou (2023). Heterogeneous effects of covid-19 vaccination on health outcomes: Evidence from administrative data. *American Economic Journal: Applied Economics*. Forthcoming.

De Chaisemartin, C. and X. d'Haultfoeuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–96.

Dean, N. E., J. W. Hogan, and M. E. Schnitzer (2021). COVID-19 vaccine effectiveness and the test-negative design. *The New England Journal of Medicine 385*(15), 1431–1433.

Deb, P., D. Furceri, J. D. Ostry, and N. Tawk (2021). The effects of covid-19 vaccination on economic activity: Evidence from us states. *IMF Working Papers 2021*(281).

Dolgin, E. et al. (2021). COVID vaccine immunity is waning—how much does that matter? *Nature 597*(7878), 606–607.

Eichner, M. J. (1998). The demand for medical care: What people pay does matter. *The American Economic Review 88*(2), 117–121.

Feldman, R. and B. Dowd (1991). A new estimate of the welfare loss of excess health insurance. *The American Economic Review 81*(1), 297–301.

Feldstein, M. S. (1973). The welfare loss of excess health insurance. *Journal of Political Economy 81*(2, Part 1), 251–280.

García-Montero, C., O. Fraile-Martínez, C. Bravo, D. Torres-Carranza, L. Sanchez-Trujillo, A. M. Gómez-Lahoz, L. G. Guijarro, N. García-Honduvilla, A. Asúnsolo, J. Bujan, et al. (2021). An updated review of sars-cov-2 vaccines and the importance of effective vaccination programs in pandemic times. *Vaccines 9*(5), 433.

Hortaçsu, A., J. Liu, and T. Schwieg (2021). Estimating the fraction of unreported infections in epidemics with a known epicenter: an application to COVID-19. *Journal of Econometrics 220*(1), 106–129.

Imbens, G. W. and J. M. Wooldridge (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature 47*(1), 5–86.

Kim, D. and Y. J. Lee (2022). Vaccination strategies and transmission of COVID-19: Evidence across advanced countries. *Journal of Health Economics 82*, 102589.

Lipsitch, M. and N. E. Dean (2020). Understanding COVID-19 vaccine efficacy. *Science 370*(6518), 763–765.

Machado, B. A. S., K. V. S. Hodel, L. M. d. S. Fonseca, V. C. Pires, L. A. B. Mascarenhas, L. P. C. da Silva Andrade, M. A. Moret, and R. Badaró (2022). The importance of vaccination in the

context of the covid-19 pandemic: a brief update regarding the use of vaccines. *Vaccines 10*(4), 591.

Manning, W. G., J. P. Newhouse, N. Duan, E. B. Keeler, and A. Leibowitz (1987). Health insurance and the demand for medical care: evidence from a randomized experiment. *The American Economic Review*, 251–277.

Mukherjee, A., K. Xu, and L. Zhang (2021). The effect of vaccination on mobility during the COVID-19 pandemic. *Nature Human Behaviour 5*(11), 1474–1483.

Okeke, E. N. (2022). Playing defense? health care in the era of COVID. *Journal of Health Economics*, 102665.

Perkins, T., M. Cavalleri, M. Lajous, and N. Ferguson (2023). Disease transmission and the economy: Evidence from the covid-19 pandemic in europe. *arXiv preprint arXiv:2211.00054*.

Perlis, R. H., K. Ognyanova, A. Uslu, K. L. Trujillo, M. Santillana, J. N. Druckman, M. A. Baum, and D. Lazer (2024). Trust in physicians and hospitals during the covid-19 pandemic in a 50-state survey of us adults. *JAMA network open 7*(7), e2424984–e2424984.

Rios-Avila, F., B. Callaway, and P. H. Sant'Anna (2022). csdid: Difference-in-differences with multiple time periods in stata. *Working paper, Boston University*.

Roth, J., P. H. Sant'Anna, A. Bilinski, and J. Poe (2022). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *arXiv preprint arXiv:2201.01194*.

Shi, M. (2024). Monitoring for waste: Evidence from medicare audits. *The Quarterly Journal of Economics 139*(2), 993–1049.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.

Tollefson, J. (2020). Why deforestation and extinctions make pandemics more likely. *Nature 584*(7820), 175–177.

Walmsley, T., A. Rose, R. John, D. Wei, J. P. Hlávka, J. Machado, and K. Byrd (2023). Macroeconomic consequences of the covid-19 pandemic. *Economic modelling 120*, 106147.

Figure 1: COVID-19 Vaccine Uptake Over Time by Age Group for U.S. Counties from Jan-2020 to Jul-2022

This figure presents monthly trends in COVID-19 vaccine uptake for two age groups—elderly (65+ years in red diamonds) and the younger adults (18-64 years in blue circles) —across U.S. counties, from January 2020 through Jul-2022. The figure is divided into three panels corresponding to: (1) population share receiving two doses of vaccine, (2) population share receiving at least one dose, and (3) population share receiving a booster dose.
These figures use census data on vaccinations obtained from CDC; this data is available daily at the county-level. As COVID-19 vaccination started in the U.S. in Dec-2020, the vaccination data is also available from this month. The following variables are available in this data: number and percentage of individuals in the county that have received one dose, two doses and the booster dose. This data is provided for the following two age groups: 18+ years and 65+ years.

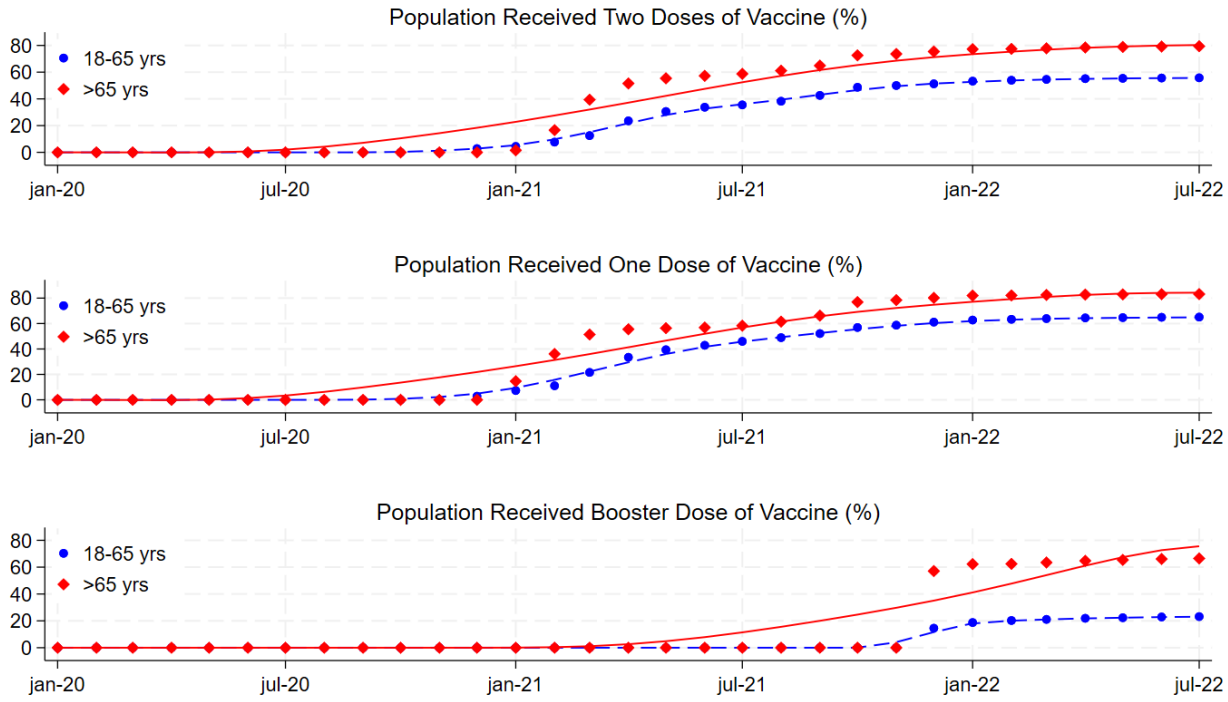


Figure 2: Geographic and Temporal Variation in Two-Dose COVID-19 Vaccine Coverage by Age Group Across U.S. Counties from Jan 2020 to Jul 2022

This figure illustrates two key dimensions of variation in COVID-19 vaccine coverage used in the identification strategy of this study: (i) the intensive margin of vaccine uptake across U.S. counties, and (ii) the extensive margin of cross-sectional heterogeneity within each month, both disaggregated by age group. The top panel shows the time series for the number of counties vaccinating at least half their elderly (65+ years in red diamonds) and young adult population (18-64 years in blue circles). The bottom panel presents monthly box plots of county-level rates of two-dose vaccination from Jan-2021 to Jul-2022, again separately for elderly (pink boxes) and younger adults (blue boxes).

These figures use census data on vaccinations obtained from CDC; this data is available daily at the county-level. As COVID-19 vaccination started in the U.S. in Dec-2020, the vaccination data is also available from this month. The following variables are available in this data: number and percentage of individuals in the county that have received one dose, two doses and the booster dose. This data is provided for the following two age groups: 18+ years and 65+ years.

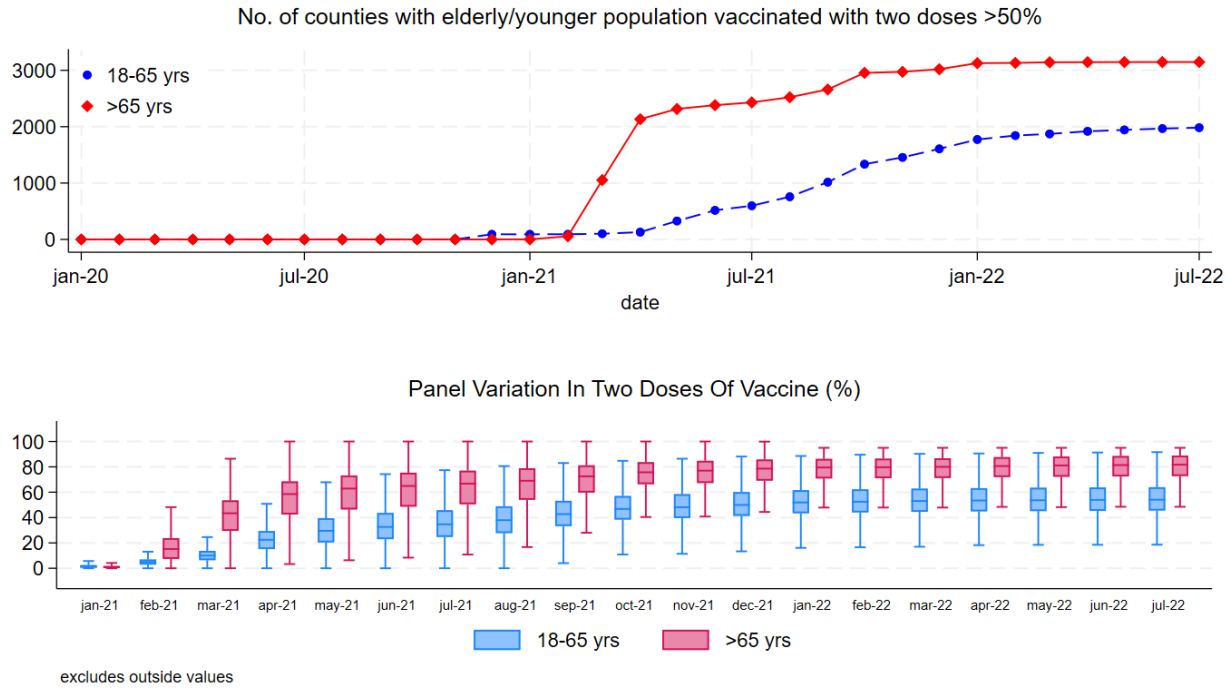


Figure 3: Trends in COVID-19 Health Outcomes by Age Group (U.S. Counties, Jan 2020–Jul 2022)

This figure presents monthly trends in four key COVID-19 health outcomes—(i) deaths per million, (ii) ICU admissions per million, (iii) hospitalizations per million, and (iv) confirmed cases per million—for two age groups: elderly individuals (65+ years, shown in red with diamonds) and younger adults (18–65 years, shown in blue with circles). The data are aggregated across U.S. counties and span from January 2020 to Jul-2022.

The figures are constructed using the census data on health outcomes from the case surveillance data on COVID-19 put out by the Centers for Disease Control and Prevention (CDC) for public use. This *anonymized* data is available at the individual level and captures whether the individual has a confirmed case of COVID-19 or not, whether the individual is symptomatic or not, whether the individual had to be hospitalized or admitted to the Intensive Care Unit (ICU) of a hospital, the date when the case was reported, the age group of the individual (0-17 years, 18-49 years, 50-64 years, and 65+ years), and the county and state the individual is located in.

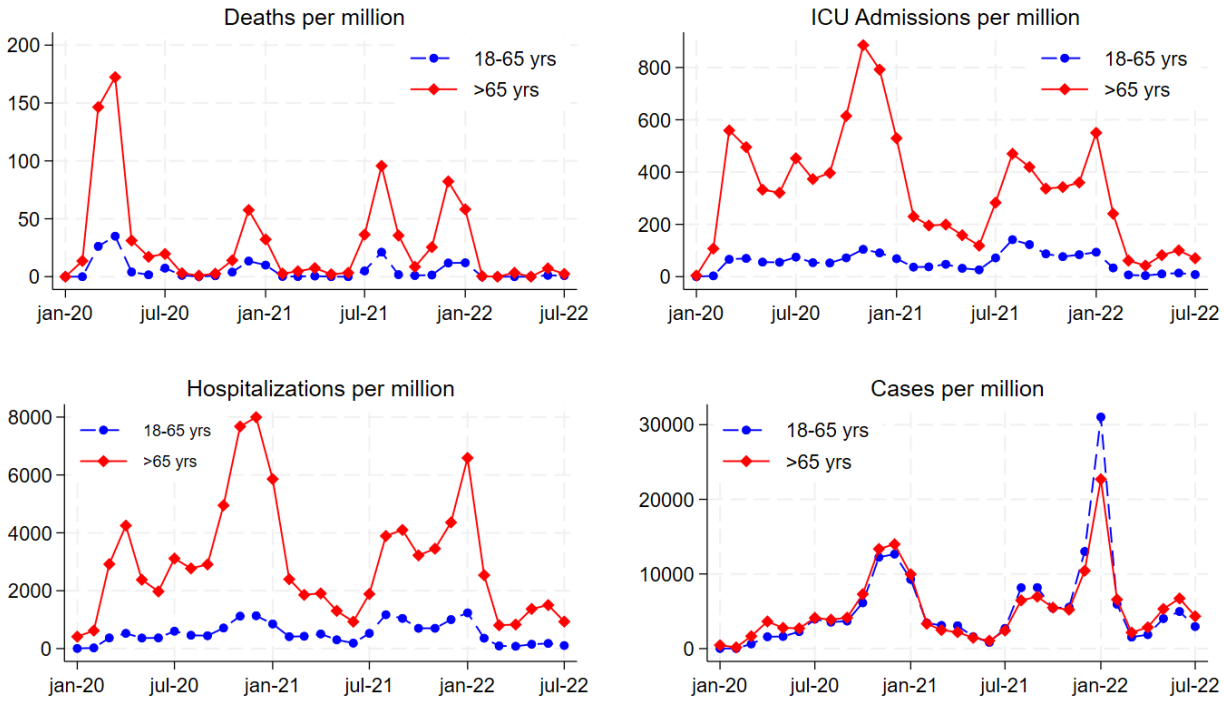


Figure 4: Pre-Vaccine COVID-19 Health Outcomes and Mar-2021 Vaccination Rates: Assessing Reverse Causality

This figure examines whether early vaccination rates in Mar-2021 were systematically correlated with the severity of local COVID-19 outcomes in Dec-2020, i.e., whether reverse causality could plausibly explain the observed variation in vaccine uptake across counties. Each panel plots the relationship between county-level COVID-19 health metrics in Dec-2020 on the x-axis and the share of each age group vaccinated by Mar-2021 on the y-axis, separately for individuals receiving at least one dose (blue circles) and two doses (red diamonds). The top panels shows the (lack of) correlation of vaccination—both one dose and two doses—for the elderly (65+ years) in Mar-2021 with deaths and cases per million three months earlier in Dec-2020. The bottom panels shows the (lack of) correlation of vaccination —both one dose and two doses—for the younger adults (18-64 years) in Mar-2021 with cases per million three months earlier in Dec-2020.

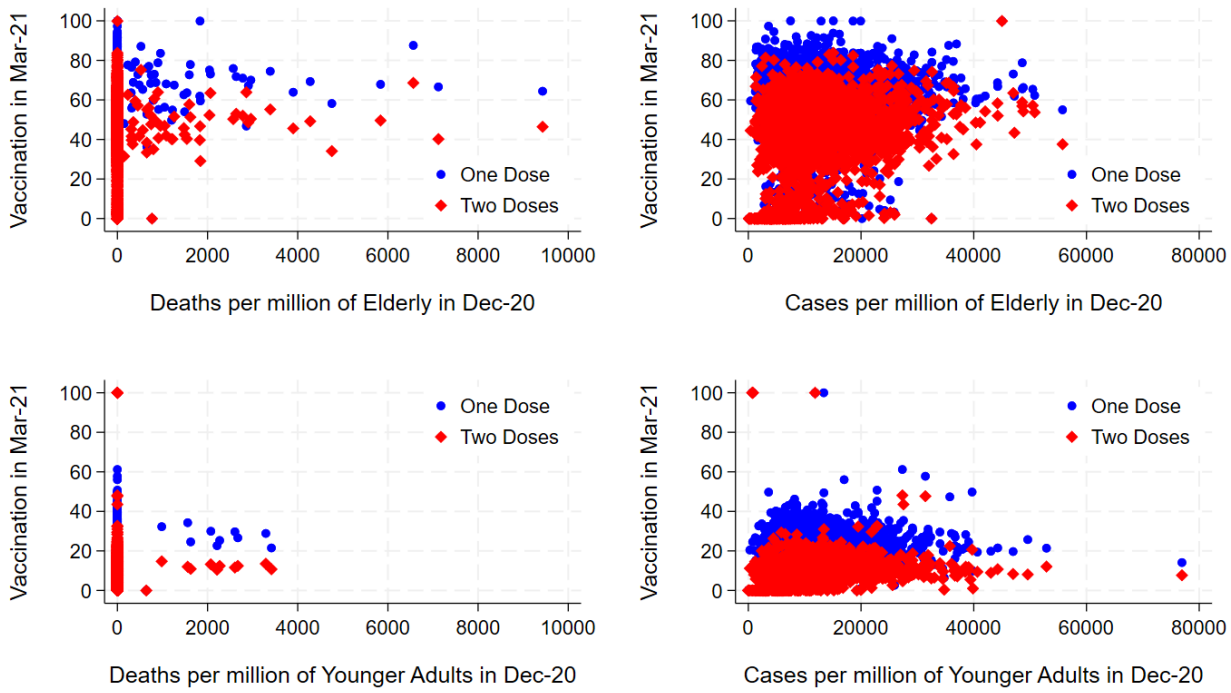
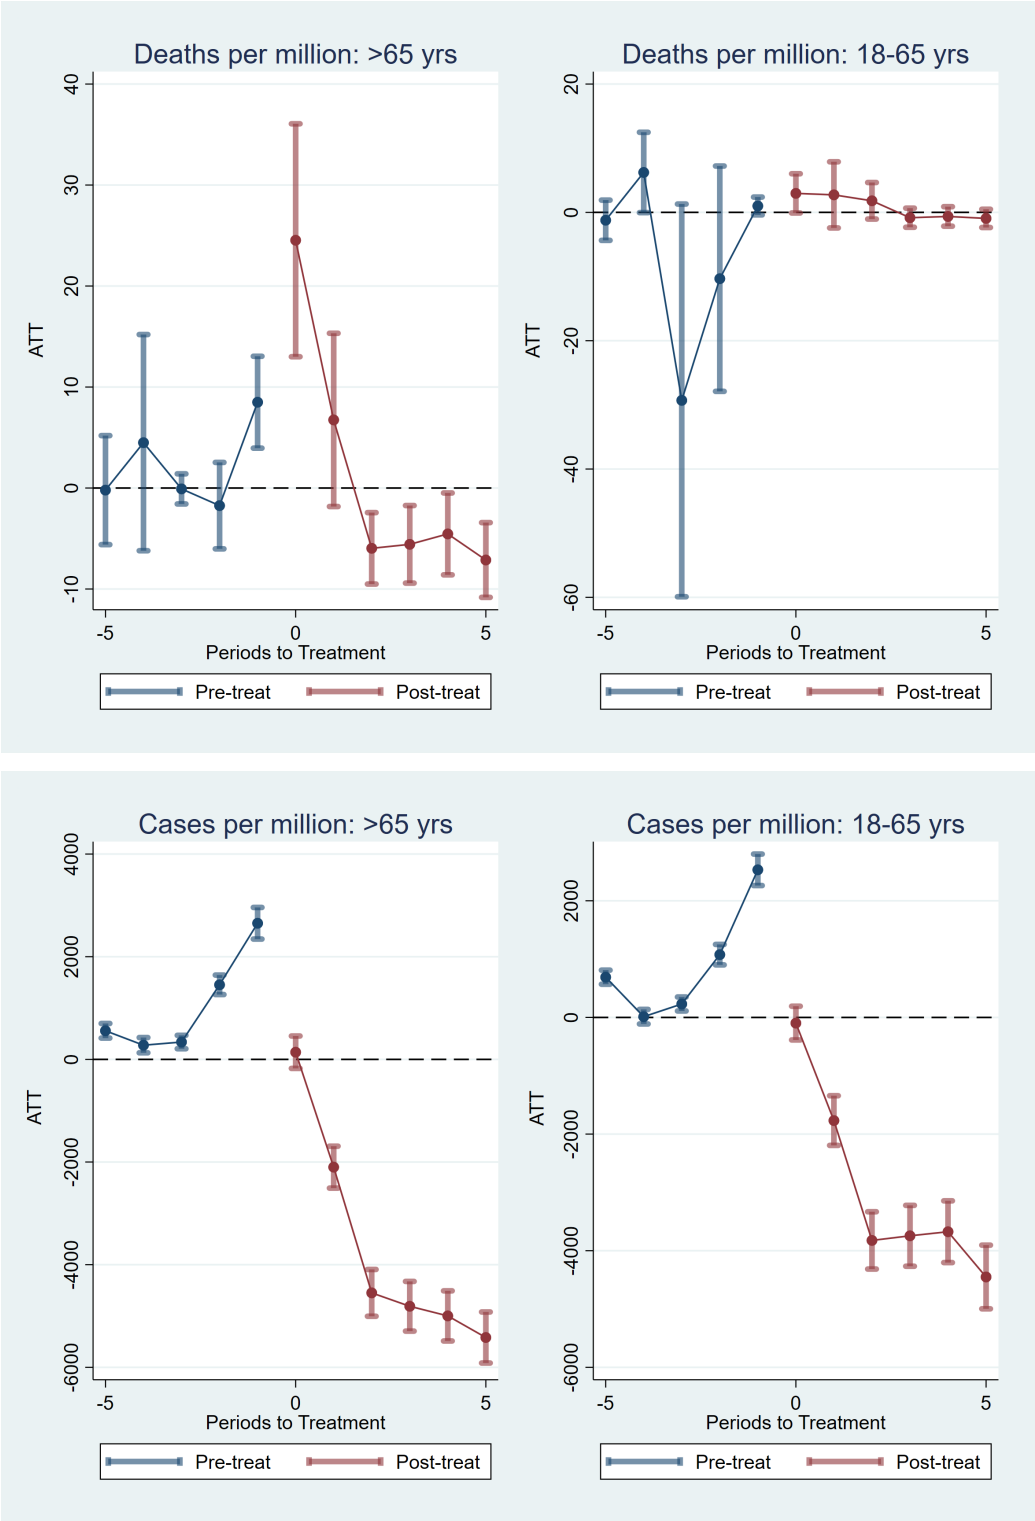
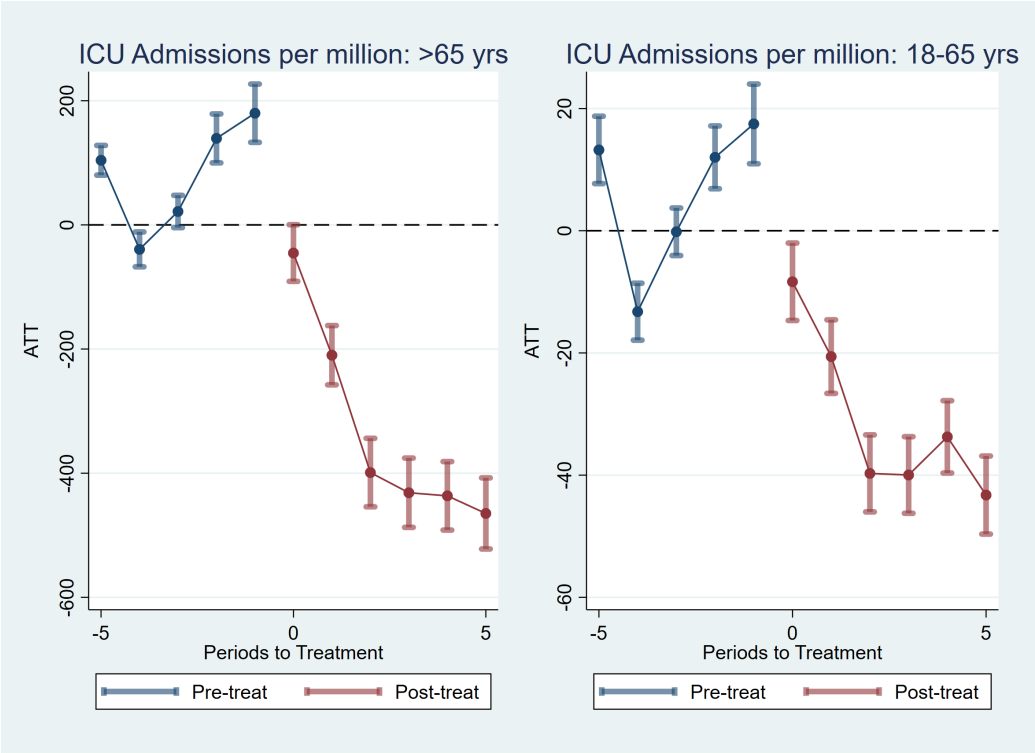


Figure 5: Dynamic effects estimated using the methodology in Rios-Avila et al. (2022) of "not yet treated" counties as control group

This figure shows the difference-in-difference estimates of COVID-19 vaccination on deaths using the Rios-Avila et al. (2022) methodology of "not yet treated" counties as control group from five months before treatment to five months after; 0 corresponds to the month of treatment when vaccination begins in a county.




ICU Admissions per million: >65 yrs
ATT
200
0
-200
-400
-600
-5
0
5
Periods to Treatment
Pre-treat
Post-treat
ICU Admissions per million: 18-65 yrs
ATT
20
0
-20
-40
-60
-5
0
5
Periods to Treatment
Pre-treat
Post-treat

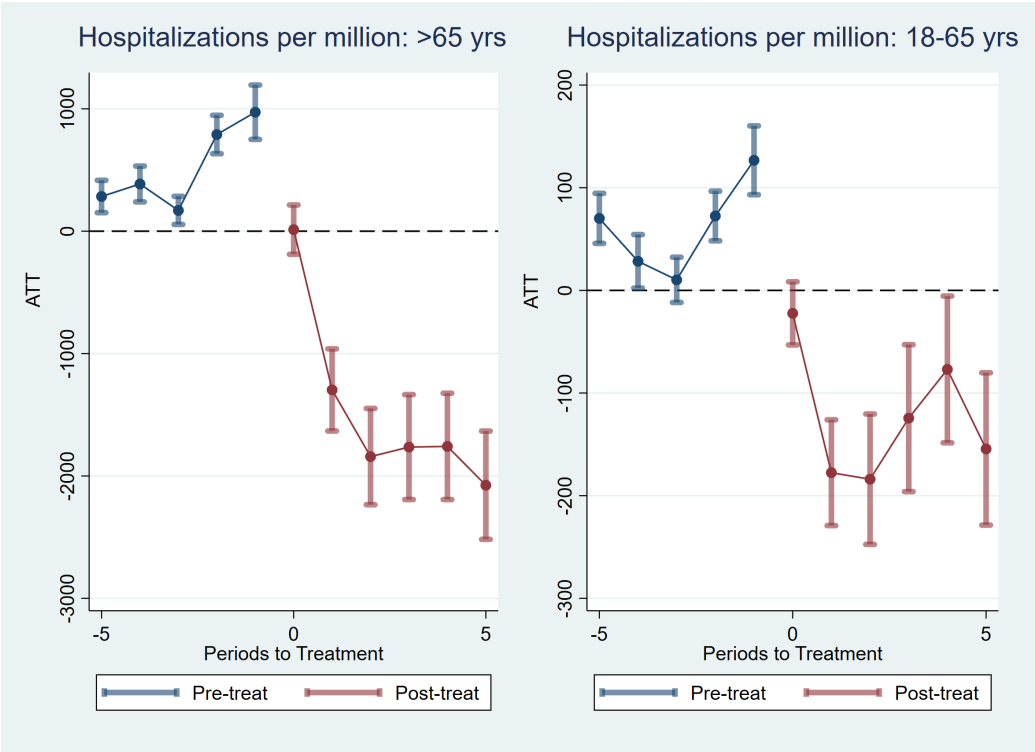
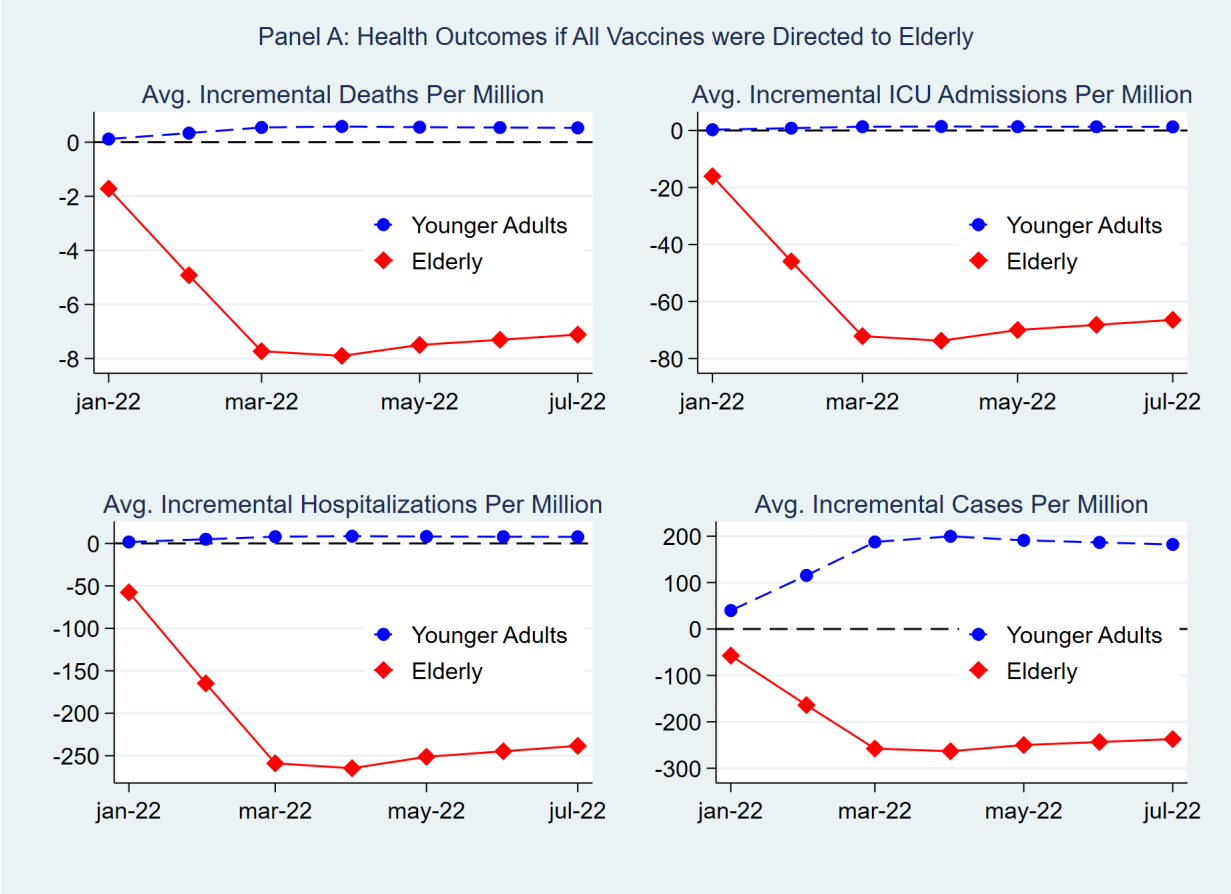
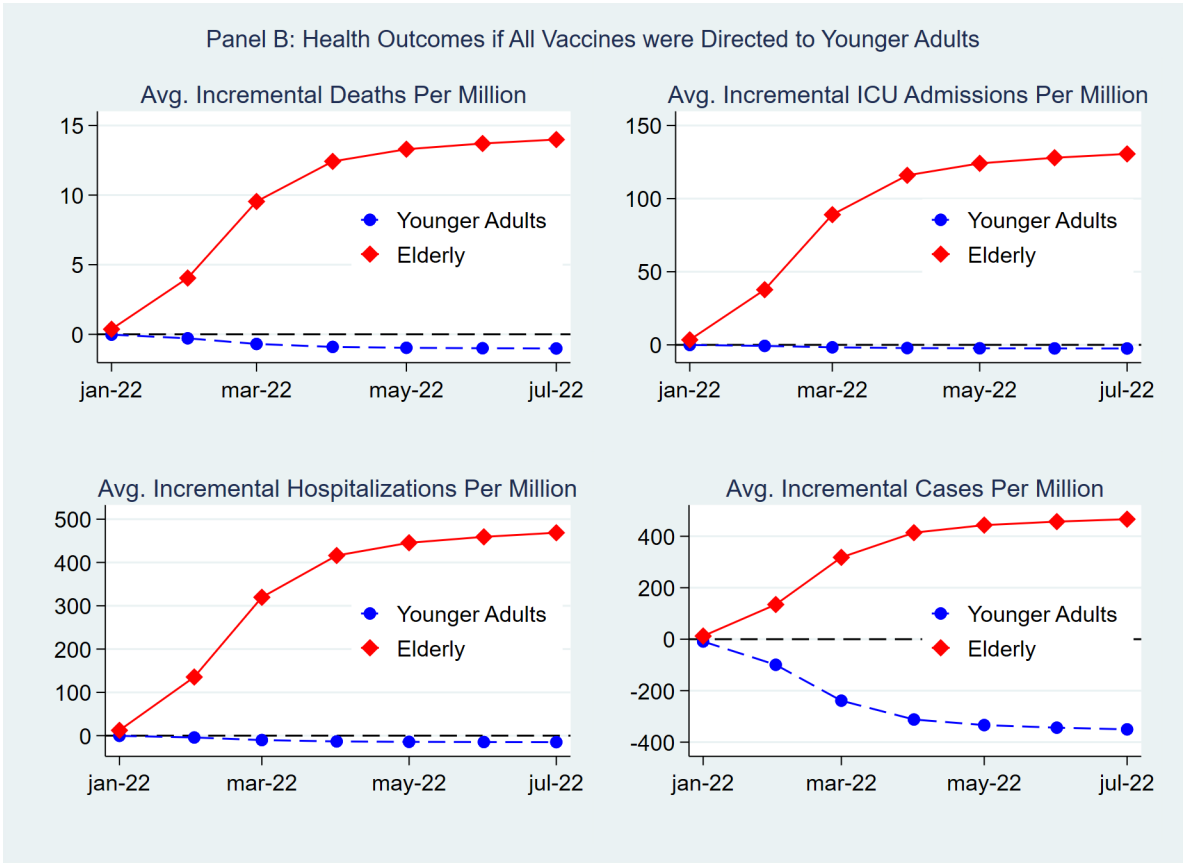

Hospitalizations per million: >65 yrs
ATT
1000
0
-1000
-2000
-3000
-5
0
5
Periods to Treatment
Pre-treat
Post-treat
Hospitalizations per million: 18-65 yrs
ATT
200
100
0
-100
-200
-300
-5
0
5
Periods to Treatment
Pre-treat
Post-treat

Figure 6: Counterfactual Estimates for Health Outcomes

This figure presents the counterfactual estimates for health outcomes for three scenarios. In each scenario, the figures show the average at the (county, month) level for *incremental* deaths per million, ICU admissions per million, hospitalizations per million, and infections per million compared to the actual health outcome that manifested. In other words, each of these measures represents the flow of incremental health outcomes in the particular month averaged across all counties. The first scenario involves directing all vaccines exclusively to the elderly (65+) until they are fully vaccinated. The second scenario involves directing all vaccines exclusively to younger adults (18–64) until they are fully vaccinated. Finally, the third scenario involves allocating vaccines proportionally to each group's share in the county's population.




Panel B: Health Outcomes if All Vaccines were Directed to Younger Adults
Avg. Incremental Deaths Per Million
Avg. Incremental ICU Admissions Per Million
Avg. Incremental Hospitalizations Per Million
Avg. Incremental Cases Per Million
Younger Adults
Elderly
jan-22
mar-22
may-22
jul-22

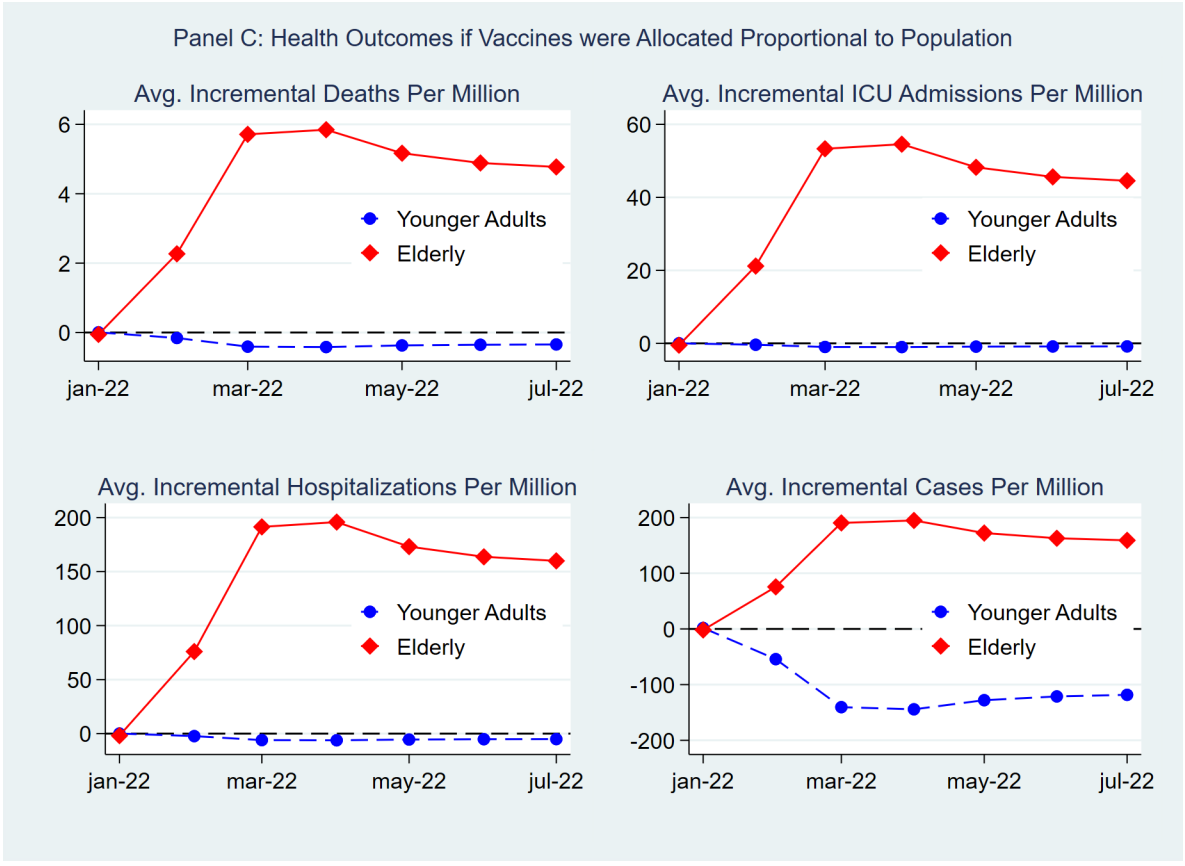

Panel C: Health Outcomes if Vaccines were Allocated Proportional to Population
Avg. Incremental Deaths Per Million
Avg. Incremental ICU Admissions Per Million
Avg. Incremental Hospitalizations Per Million
Avg. Incremental Cases Per Million
Younger Adults
Elderly
jan-22
mar-22
may-22
jul-22

Table 1: Summary Statistics of the Sample

This table displays the summary statistics of the sample. Panels A and B show for the elderly and for younger adults respectively.

Variable:	(1) N	(2) Mean	(3) Std. devn.	(4) Min.	(5) Max.
Panel A: Elderly					
Deaths per million	99,975	10.7	192.8	0	14,410.8
ICU Admissions per million	99,975	142.0	683.3	0	29,255.3
Hospitalizations per million	99,975	1,327.9	3,238.9	0	76,288.7
Cases per million	99,975	2,534.0	5,457.7	0	166,209.4
Population received two doses (%)	99,975	38.0	37.5	0	100.0
Population received booster dose (%)	99,975	16.4	28.8	0	95.0
Panel B: Younger Adults					
Deaths per million	99,975	2.5	84.4	0	8,213.8
ICU Admissions per million	99,975	28.2	133.7	0	3,908.9
Hospitalizations per million	99,975	271.0	668.1	0	18,213.5
Cases per million	99,975	2,791.9	6,215.7	0	118,019.9
Population received two doses (%)	99,975	24.7	27.1	0	100.0
Population received booster dose (%)	99,975	5.2	10.3	0	85.2

Table 2: Effect Of COVID-19 Vaccine On Health Outcomes - Triple-Difference Tests

This table displays the results of panel regressions examining the impact of COVID-19 vaccination on deaths, ICU admissions, hospitalizations, symptomatic cases, and confirmed cases, all per million of population. The OLS regression implemented is:

$y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn_{ic(t-12)} \cdot \Lambda_{65+,i} + \beta_2 Vaccn_{ic(t-12)} + \beta_3 \cdot \Lambda_{65+,i} + \varepsilon_{ct}$, and

where the variable $Vaccn_{ic(t-12)}$ equals the cumulative proportion of the population in age group i that has received two doses of the COVID-19 vaccine in county c until one year earlier, i.e., by month $t-12$. The dummy variable $\Lambda_{65+,i}$ equals 1 for population above 65 years of age and 0 otherwise. β_c and β_t denote county and time fixed effects respectively. Monthly data from January 2020 to Jul-2022 are obtained from the CDC. Robust standard errors clustered at the county-level are shown in parentheses. ***, **, and * denote statistical significance at the 99%, 95% and 90% levels respectively.

Dependent variable per million:	(1) Deaths	(2) ICU	(3) Hospital	(4) Cases
Population received two doses 1-year before (%) *	-0.10***	-1.50***	-10.47***	19.71***
Dummy for 65+ years	(0.01)	(0.09)	(0.46)	(0.93)
Population received two doses 1-year before (%)	-0.12***	-0.11	1.48***	-21.28***
	(0.03)	(0.09)	(0.56)	(1.69)
Dummy for 65+ years	9.56***	127.80***	1,145.56***	-347.25***
	(0.88)	(5.94)	(26.89)	(15.18)
Observations	199,950	199,950	199,950	199,950
Adjusted R-squared	0.06	0.22	0.31	0.43
County FE & Time FE	Yes	Yes	Yes	Yes

Table 3: Effect of Two Doses of COVID-19 Vaccine on Health Outcomes of Elderly and Younger Adults - Difference-in-difference Tests

This table displays the results of panel regressions examining the impact of two doses of COVID-19 vaccine on deaths, ICU admissions, hospitalizations, symptomatic cases and confirmed cases, all per million of population. The OLS regression implemented is:

$y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn_{ic(t-12)} + \varepsilon_{ct}$,

where β_c and β_t denote the fixed effects for county and sample month respectively, and all the other variables are as defined previously. Monthly data from January 2020 to Jul-2022 are obtained from the CDC. Robust standard errors clustered at the county-level are shown in parentheses. ***, **, and * denote statistical significance at the 99%, 95% and 90% levels respectively.

Dependent variable per million:	(1) Deaths	(2) ICU	(3) Hospital	(4) Cases
Panel A: Elderly				
Population received two doses 1-year before (%)	-0.24***	-2.24***	-8.04***	-8.00***
	(0.05)	(0.14)	(0.85)	(1.46)
Observations	99,975	99,975	99,975	99,975
Adjusted R-squared	0.07	0.30	0.39	0.42
County FE & Time FE	Yes	Yes	Yes	Yes
Panel B: Younger Adults				
Population received two doses 1-year before (%)	-0.07***	-0.17***	-1.03***	-24.11***
	(0.02)	(0.03)	(0.23)	(1.83)
Observations	99,975	99,975	99,975	99,975
Adjusted R-squared	0.06	0.31	0.40	0.44
County FE & Time FE	Yes	Yes	Yes	Yes

Table 4: Dynamic Effects of COVID-19 Vaccination with Varying Lags

This table presents triple-difference (Panel A) and age-group-specific difference-in-differences (Panels B and C) estimates for the effect of COVID-19 vaccination on health outcomes, using lags ranging from contemporaneous (same month) to one year prior. For brevity, I display **only the coefficient estimate** and statistical significance for the triple-difference estimates in Panel A, difference-in-difference estimates for the elderly in Panel B and the same for younger adults in Panel C—all represented by β_1—in the regression equations below. Panel A implements the following OLS regression for the triple-difference estimate:

$$y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn_{ic(t-m)} \cdot \Lambda_{65+,i} + \beta_2 Vaccn_{ic(t-m)} + \beta_3 \cdot \Lambda_{65+,i} + \varepsilon_{ct},\ m = 0, 3, 6, 12$$

Panels B and C implement the following OLS regression for the difference-in-difference estimates:

$$y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn_{ic(t-m)} + \varepsilon_{ct},\ m = 0, 3, 6, 12$$

Monthly data from January 2020 to Jul-2022 are obtained from the CDC. Robust standard errors clustered at the county-level are shown in parentheses. ***, **, and * denote statistical significance at the 99%, 95% and 90% levels respectively.

	(1)	(2)	(3)	(4)
Dependent variable per million:	Deaths	ICU	Hospital	Cases
Panel A: Triple-difference using different lags in vaccination				
2-dose vaccination **one year before**	-0.16***	-1.59***	-10.51***	12.22***
2-dose vaccination six months before	0.23***	1.53***	16.32***	-11.49***
2-dose vaccination three months before	0.24***	1.53***	16.13***	-15.39***
2-dose vaccination in the same month	0.21***	1.57***	16.08***	-14.66***
County FE & Time FE	Yes	Yes	Yes	Yes
Panel B: Difference-in-difference (Elderly) using different lags in vaccination				
2-dose vaccination **one year before**	-0.24***	-2.24***	-8.04***	-8.00***
2-dose vaccination six months before	-0.24**	0.08	2.88***	0.11
2-dose vaccination three months before	-0.03	0.09	4.13***	8.76***
2-dose vaccination in the same month	0.14**	-0.64***	-0.19	15.14***
County FE & Time FE	Yes	Yes	Yes	Yes
Panel C: Difference-in-difference (Younger Adults) with different lags in vaccination				
2-dose vaccination **one year before**	-0.07***	-0.17***	-1.03***	-24.11***
2-dose vaccination six months before	-0.17***	-0.11***	-0.28	20.08***
2-dose vaccination three months before	-0.13***	0.06*	0.54**	29.33***
2-dose vaccination in the same month	-0.08	-0.03	-0.51**	26.44***
County FE & Time FE	Yes	Yes	Yes	Yes

Table 5: Incremental Effect Of Two Doses of COVID-19 Vaccine One Year Before vis-à-vis Booster Dose Now On Health Outcomes - Triple-Difference Tests

This table displays the results of the following panel regressions:

$y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn2Doses_{ic(t-12)} \cdot \Lambda_{65+,i} + \beta_2 VaccnBooster_{c,t} \cdot \Lambda_{65+,i} + \beta_1 Vaccn2Doses_{ic(t-12)} + \beta_4 \Lambda_{65+,i} + \varepsilon_{ct}$,

where the variable $VaccnBooster_{ct}$ equal the cumulative proportion of the population in age group i that has received the booster dose of the COVID-19 vaccine respectively in county c by month t. The dummy variable $\Lambda_{65+,i}$ equals 1 for population above 65 years of age and 0 otherwise. β_c and β_c denote county and time fixed effects respectively. Monthly data from January 2020 to Jul-2022 are obtained from the CDC. Robust standard errors clustered at the county-level are shown in parentheses. ***, **, and * denote statistical significance at the 99%, 95% and 90% levels respectively.

	(1)	(2)	(3)	(4)
Dependent variable per million:	Deaths	ICU	Hospital	Cases
Population received two doses 1-year before (%) *	-0.38***	-2.05***	-22.05***	68.98***
Dummy for 65+ years	(0.04)	(0.16)	(0.93)	(5.61)
Population received booster dose now (%) *	0.33***	0.82***	13.67***	-150.72***
Dummy for 65+ years	(0.04)	(0.16)	(0.88)	(6.29)
Population received two doses 1-year before (%)	-0.02	-0.51***	-3.99***	-141.32***
	(0.02)	(0.09)	(0.70)	(6.14)
Population received booster dose now (%)	-0.13***	-0.40***	-2.42***	214.65***
	(0.04)	(0.12)	(0.71)	(7.07)
Dummy for 65+ years	7.85***	125.48***	1,076.20***	-221.89***
	(0.79)	(5.80)	(25.39)	(14.20)
Observations	199,950	199,950	199,950	199,950
Adjusted R-squared	0.06	0.20	0.25	0.25
County FE & Time FE	Yes	Yes	Yes	Yes

Table 6: Robustness of Results when Data is Aggregated to MSA Level

This table displays the results of panel regressions examining the impact of COVID-19 vaccine on deaths, ICU admissions, hospitalizations and confirmed cases per million by aggregating the data to the MSA level. Panel A implements the following OLS regression for the triple-difference estimate:

$y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn_{ic(t-12)} \cdot \Lambda_{65+,i} + \beta_2 Vaccn_{ic(t-12)} + \beta_3 \cdot \Lambda_{65+,i} + \varepsilon_{ct}$

Panels B and C implement the following OLS regression for the difference-in-difference estimates:

$y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn_{ic(t-12)} + \varepsilon_{ct}$

Monthly data from January 2020 to Jul-2022 are obtained from the CDC. Robust standard errors clustered at the MSA-level are shown in parentheses. ***, **, and * denote statistical significance at the 99%, 95% and 90% levels respectively.

Dependent variable per million:	(1) Deaths	(2) ICU	(3) Hospital	(4) Cases
Panel A: Triple-difference estimates				
Population received two doses 1-year before (%) *	-0.57***	-2.87***	-20.29***	16.46***
Dummy for 65+ years	(0.07)	(0.31)	(1.55)	(3.54)
Population received two doses 1-year before (%)	-0.15**	-0.82***	-5.31***	-24.30***
	(0.06)	(0.20)	(1.43)	(4.12)
Dummy for 65+ years	37.59***	248.33***	2,260.46***	-620.86***
	(4.26)	(20.86)	(70.89)	(47.09)
Observations	26,286	26,286	26,286	26,286
Adjusted R-squared	0.08	0.29	0.26	0.05
MSA FE & Time FE	Yes	Yes	Yes	Yes
Panel B: Difference-in-difference estimates for the Elderly				
Population received two doses 1-year before (%)	-0.77***	-2.05***	-2.04	-0.27
	(0.28)	(0.39)	(2.37)	(4.47)
Observations	13,143	13,143	13,143	13,143
Adjusted R-squared	0.10	0.43	0.47	0.54
MSA FE & Time FE	Yes	Yes	Yes	Yes
Panel C: Difference-in-difference estimates for the Younger Adults				
Population received two doses 1-year before (%)	-0.15*	-0.06	0.07	-35.72***
	(0.08)	(0.10)	(0.59)	(6.50)
Observations	13,143	13,143	13,143	13,143
Adjusted R-squared	0.10	0.42	0.45	0.58
MSA FE & Time FE	Yes	Yes	Yes	Yes

Table 7: Robustness of Results when different fixed effects are employed for analysis

This table displays the results of panel regressions examining the impact of COVID-19 vaccine on deaths, ICU admissions, hospitalizations and confirmed cases per million by aggregating the data to the county level. The set of fixed effects employed are different here from those used in Tables 2. Monthly data from January 2020 to Jul-2022 are used; all the data are obtained from the CDC. Robust standard errors clustered at the county-level are shown in parentheses. ***, **, and * denote statistical significance at the 99%, 95% and 90% levels respectively.

	(1)	(2)	(3)	(4)
Dependent variable per million:	Deaths	ICU	Hospital	Cases
Panel A: Triple-difference using FE for each (County, Time) pair				
Population received two doses 1-year before (%) *	-0.16***	-1.59***	-10.51***	12.22***
Dummy for 65+ years	(0.02)	(0.13)	(0.72)	(0.77)
Population received two doses 1-year before (%)	0.03***	0.45***	2.20**	1.33
	(0.01)	(0.11)	(1.08)	(0.87)
Dummy for 65+ years	9.57***	126.24***	1,142.84***	-374.07***
	(1.23)	(8.28)	(37.77)	(21.07)
Observations	199,950	199,950	199,950	199,950
Adjusted R-squared	0.50	0.26	0.37	0.91
(County*Time) FE	Yes	Yes	Yes	Yes
Panel B: Triple-difference using (County*Age Group) and Time FE				
Population received two doses 1-year before (%) ×	-0.12***	-1.75***	-11.35***	23.04***
Dummy for 65+ yrs	(0.02)	(0.09)	(0.36)	(1.04)
Population received two doses 1-year before (%)	-0.12***	-0.44***	-1.04**	-24.34***
	(0.03)	(0.06)	(0.43)	(1.85)
Dummy for 65+ yrs EXCLUDED as (County*Age Group) FE subsume its effect				
Observations	199,950	199,950	199,950	199,950
Adjusted R-squared	0.07	0.30	0.39	0.42
(County*Age Group) FE and Time FE	Yes	Yes	Yes	Yes
Panel C: Triple-difference using State-specific time trends as well as County and Time FE				
Population received two doses 1-year before (%) ×	-0.16***	-1.40***	-9.65***	19.51***
Dummy for 65+ yrs	(0.02)	(0.10)	(0.57)	(1.02)
Population received two doses 1-year before (%)	0.05	-0.31**	-0.91	-19.35***
	(0.05)	(0.12)	(0.87)	(1.96)
Dummy for 65+ yrs	9.51***	128.13***	1,151.00***	-353.40***
	(0.87)	(5.98)	(27.02)	(14.83)
Observations	199,950	199,950	199,950	199,950
Adjusted R-squared	0.07	0.22	0.31	0.45
State-specific Time Trends	Yes	Yes	Yes	Yes
County and Time FE	Yes	Yes	Yes	Yes

Table 8: Robustness of Results when data only for 2021 and 2022 are used for analysis

This table displays the results of panel regressions examining the impact of COVID-19 vaccine on deaths, ICU admissions, hospitalizations and confirmed cases per million by aggregating the data to the county level. Data for 2020 is excluded from these tests to examine if the main results are driven by the presence of zeros for vaccination in 2020. Thus, monthly data from Jan-2021 to Jul-2022 are used; all the data are obtained from the CDC. Robust standard errors clustered at the county-level are shown in parentheses. ***, **, and * denote statistical significance at the 99%, 95% and 90% levels respectively.

	(1)	(2)	(3)	(4)
Dependent variable per million:	Deaths	ICU	Hospital	Cases
Panel A: Triple-difference estimates				
Population received two doses 1-year before (%) *	-0.18***	-1.10***	-8.96***	20.04***
Dummy for 65+ years	(0.02)	(0.09)	(0.53)	(1.40)
Population received two doses 1-year before (%)	-0.05***	-0.82***	-6.37***	-35.88***
	(0.01)	(0.07)	(0.49)	(1.64)
Dummy for 65+ years	11.73***	113.64***	1,149.78***	-510.59***
	(1.11)	(5.92)	(28.41)	(21.16)
Observations	122,550	122,550	122,550	122,550
Adjusted R-squared	0.06	0.26	0.31	0.23
County FE & Time FE	Yes	Yes	Yes	Yes
Panel B: Difference-in-difference estimates for the Elderly				
Population received two doses 1-year before (%)	-0.18***	-1.99***	-7.37***	-3.21**
	(0.05)	(0.15)	(0.88)	(1.41)
Observations	61,275	61,275	61,275	61,275
Adjusted R-squared	0.07	0.36	0.46	0.44
County FE & Time FE	Yes	Yes	Yes	Yes
Panel C: Difference-in-difference estimates for the Younger Adults				
Population received two doses 1-year before (%)	-0.03***	-0.18***	-1.18***	-28.98***
	(0.01)	(0.04)	(0.24)	(1.89)
Observations	61,275	61,275	61,275	61,275
Adjusted R-squared	0.07	0.35	0.43	0.45
County FE & Time FE	Yes	Yes	Yes	Yes

Table 9: Validation and Robustness of Triple-Difference and Difference-in-difference Estimates Across Different Specifications

This table has three parts. Panel A summarizes the core triple-difference coefficients across a battery of robustness checks, including alternative fixed effects structures (e.g., county-time FE, county-by-age and time FE, state-specific time trends added to county and time FE), analysis at the MSA level, and models controlling for booster doses. The consistent sign, magnitude, and statistical significance of the coefficients across these diverse specifications demonstrate the robustness of the main findings to a wide range of potential confounding factors. Panels B and C provide external validation by comparing the baseline triple-difference estimates with the results from separate difference-in-differences (DiD) analyses for the elderly and younger adults, respectively. For brevity, **only the coefficient** and statistical significance of the key estimates are shown. ***, **, and * denote statistical significance at the 99%, 95% and 90% levels respectively.

Panel A: Coefficients of Triple-difference Tests

Dependent variable per million:	(1) Deaths	(2) ICU	(3) Hospital	(4) Cases
Baseline (county and time FE) in Table 2	-0.10***	-1.50***	-10.47***	19.71***
(County, time) FE in Panel A of Table 7	-0.16***	-1.59***	-10.51***	12.22***
(County*Age Group) FE and Time FE from Panel B of Table 7	-0.12***	-1.75***	-11.35***	23.04***
State-specific time trends + county FE and time FE from Panel C of Table 7	-0.16***	-1.40***	-9.65***	19.51***
Controlling for booster dose from Table 5	-0.38***	-2.05***	-22.05***	68.98***
MSA level tests from Table 6	-0.57***	-2.87***	-20.29***	16.46***
Excluding observations in year 2020 from Table 8	-0.18***	-1.10***	-8.96***	20.04***

Panel B: Coefficients of Difference-in-difference Tests for Elderly

Original DiD estimates from Table 3	-0.24***	-2.24***	-8.04***	-8.00***
Triple-difference tests Table 2	-0.22***	-1.50***	-8.99***	-1.57***

Panel C: Coefficients of Difference-in-difference Tests for Younger Adults

Original DiD estimates from Table 3	-0.07***	-0.17***	-1.03***	-24.11***
Triple-difference tests from Table 2	-0.12***	-0.11	1.48***	-21.28***