Designing Pandemic Vaccination: Private and Public Returns in Optimal Allocation

Rishabh Subramanian*[1]

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

October 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 transmission by targeting socially active groups. This paper provides population-level evidence to quantify both margins for the elderly (65+) and younger adults (18–64)—evidence that cannot be obtained from clinical trials. Using monthly data for all U.S. counties from January 2020 to July 2022, I exploit the staggered vaccine roll-out and differential timing across age groups, combining a triple-difference design with difference-in-differences estimates for each group. A 10% rise in vaccination reduces deaths, ICU admissions, and hospitalizations by 3-, 13-, and 8-fold more among older adults than younger adults, respectively, while infections fall 3-fold more among younger adults. I document significant inter-generational spillovers: vaccinating younger adults significantly reduces severe outcomes among the elderly, while protecting the elderly generates indirect benefits by reducing transmission pressure. Emphasizing the long-run impact of vaccines, I measure all effects one year after vaccination. Counterfactual simulations incorporating these spillovers show that accounting for externalities reshapes predicted outcomes but does not overturn the advantage of prioritizing the elderly, whose smaller population share enables rapid coverage and large health gains. Together, these results provide quantitative evidence on how clinical vulnerability, demographic structure, and spillovers jointly influence vaccine allocation in 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 therefore: 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. Clini-

[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

cal trials and cohort studies cannot provide such evidence: they identify individual-level efficacy but cannot measure population spillovers or the general-equilibrium effects of reallocating doses across groups. Only population-level analysis can capture the full interplay between biology, behavior, and transmission dynamics.

This paper fills that gap. Using monthly, population-level data from over 3,000 U.S. counties from January 2020 to July 2022, I combine a triple-difference design exploiting age-based vaccine prioritization with difference-in-difference estimates that use variation in county-level vaccine uptake over time. The TD design isolates differences in vaccine returns across groups under shared local conditions, while the DiD estimates quantify effects within each group. Together, they yield causal estimates of the long-run, population-level health effects of vaccination. Critically, the TD framework allows me to directly test the core mechanism that the returns to vaccination vary with epidemiological risk. Because the elderly are biologically more vulnerable to death and severe outcomes, they derive larger private protective benefits from vaccination; because younger adults are more central to transmission, they generate greater public spillovers when vaccinated.

These mechanisms are reflected in the outcomes analyzed in the paper, which separately capture the private and public dimensions of vaccine benefits. Deaths, ICU admissions, and hospitalizations serve as proxies for the private returns to 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]

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. 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 obtained over shorter horizons.

As further evidence on the underlying mechanism, I find that vaccination reduces deaths per case and ICU admissions per case far more among the elderly than among younger adults. This confirms that the larger health gains for the elderly reflect true biological protection—not just differential exposure—and reinforces the view that clinical vulnerability drives the private returns from vaccination.

I then quantify the *spillover effects* of vaccinating one age group on the health outcomes of the other. The analysis shows that vaccinating younger adults significantly improves outcomes for the

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

elderly—reducing deaths and infections—while vaccinating the elderly generates smaller but still meaningful indirect benefits for younger adults. These inter-generational externalities are large enough to reshape predicted health outcomes and are critical for understanding the full health implications of vaccination policy.

To translate these findings into actionable insights for policymakers, I *simulate counterfactual scenarios* under three vaccine allocation strategies: prioritizing only the elderly, prioritizing only younger adults, and distributing doses proportional to population shares. In each case, I do not change the total vaccine supply in a county, thereby only redistributing it differently depending upon the vaccination strategy. Motivated by the evidence on inter-generational spillovers, each strategy is evaluated under two scenarios: one that incorporates spillovers across age groups and another that considers only direct within-group effects. Three key insights emerge. First, incorporating spillovers substantially reshapes predicted outcomes, especially for younger adults, whose simulated health outcomes flip from worsening to improving once spillovers are accounted for. Second, despite these strong cross-group externalities, the elderly-first strategy remains optimal, driven by the dual forces of clinical vulnerability and demographic leverage: because the elderly face higher mortality risk and constitute a smaller population share, reallocating doses toward them rapidly achieves large health gains.

In sum, this paper provides the first population-level causal evidence that jointly quantifies the private protective returns and the public-good spillovers from vaccination. By linking these heterogeneous effects to counterfactual policy simulations, the analysis clarifies how biological vulnerability, demographic structure, and social externalities jointly determine the welfare effects of vaccination policy. These findings extend beyond COVID-19, offering a framework for designing vaccine allocation strategies in future 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). While these contributions

underscore the importance of targeting, they 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. As this duality lies at the core of pandemic policy design, I extend this perspective by quantifying both protective and transmission-reducing returns to COVID-19 vaccination.

In particular, Almond et al. (2010) demonstrate that marginal returns to medical care are highly heterogeneous, with the largest gains accruing to the most vulnerable sections of the population. This mirrors the elevated returns from vaccinating the elderly observed in this study. Hendren and Sprung-Keyser (2020) develop a unified welfare framework that evaluates government policies by combining heterogeneous treatment effects, allowing policymakers to compare allocations not only by average impacts but also by their distribution of gains and losses. This approach directly parallels the allocation problem examined here: choosing between protecting the clinically vulnerable and curbing spread through socially active groups.

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: 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 I employ 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+ yrs) 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 (18–64), 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, which is available daily at the county level, provide 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. Using this data, I construct the age groups for the analysis: the elderly (65+ yrs) and younger adults (18–64 yrs). As COVID-19 vaccination started in the U.S. in Dec-2020, the vaccination data is also available from this month.

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 to Jul-2022. As the vaccination for both the elderly (65+ years) and the younger adults plateaued by early 2022, I end the sample in 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 18–64 age group, reflecting the age-based prioritization described above. The top panel shows the cumulative share of each population that completed the 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 18–64 group reached only about 55% coverage for two doses over the same period.

The middle panel, which tracks the population receiving at least one dose of the vaccine, 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 18–64 group remains below 25%.

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 I exploit to estimate the heterogeneous treatment effects.

3.3 Cross-Sectional and Temporal Variation Relevant for Identification

Figure 2 complements the insights obtained from these trends. 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 18–64). We observe here as well that elderly vaccination surged early in 2021, with over 2,000 counties surpassing 50% coverage by Mar-2021. In contrast, uptake among the 18–64 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.

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 18–64 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 provided 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). 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), where the second category is calculated using the data for 18-49, 50-64, and 65+ categories in the CDC data. 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 starts in Jan-2020 and ends in Jul-2022 and contains 199,950 observations (31 months * 2 age groups * 3,225 counties).

4.2 Proxies for Health Outcomes

I employ the following proxies for the health outcomes related to COVID-19, distinguishing between measures that capture *public returns* and those that capture *private returns* to vaccination. To measure public returns, I use total confirmed cases per million population, which reflect how vaccination reduces community transmission and thereby benefits even those not directly vaccinated; all the results are quantitatively similar when I include unconfirmed cases. To measure private returns, I examine deaths per million population—my most salient proxy for individual protection against severe outcomes—along with ICU admissions and hospitalizations per million. These outcomes capture the direct health benefits of vaccination for those at highest clinical risk. All health outcomes are measured as flow variables, i.e., the incidence of a given health event 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: Private and Public Health Burdens across Age Groups

Table 1 displays the summary statistics and Figure 3 plots the monthly averages of the health outcomes from Jan-2020 to Jul-2022. As the data is organized at the (county, month) level, the

[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.

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. The measures capturing private health burdens—deaths, ICU admissions, and hospitalizations per million—and the measure capturing public health burdens—cases per million—all show distinct peaks in Apr-2020, Dec-2020, Aug-2021, and Jan-2022. The first peak for cases per million is less visually pronounced but appears in the data as a brief spike 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 younger adults. In each subsequent wave, deaths have been significantly lower, even though other outcomes such as cases, ICU admissions, and hospitalizations were far greater.

3. The Omicron wave (Dec 2021–Jan 2022) was the most intense in terms of cases, ICU admissions, and hospitalizations, all proxies for either public burdens (cases) or private burdens (severe outcomes). Yet deaths per million were lower in this wave than in any of the others, consistent with both changes in the virus and greater vaccination coverage.

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 (142 versus 28), and hospitalizations are 4.9 times higher (1328 versus 271). These disparities illustrate the heightened private health risk borne by the elderly due to age-related vulnerabilities and comorbidities.

5. *By contrast, total cases per million—the primary proxy for public spillovers—are almost identical across the two groups.* Figure 3 shows overlapping curves for elderly and younger adults throughout most of the sample. Table 1 indicates that cases per million were only 9% lower among the elderly (2534 versus 2792). Infection rates tracked closely during national surges, with the largest divergence occurring during the Omicron wave, when younger adults' case rates exceeded those of the elderly, consistent with higher transmissibility and behavioral factors.

These descriptive patterns illustrate the *heterogeneous impact of COVID-19 across age groups in terms of private and public returns.* The elderly consistently face much larger private health burdens—deaths, ICU admissions, and hospitalizations—while the younger adults seem to play a key role in transmission. This underscores the potential for age-targeted vaccination strategies to yield differentiated returns: vaccinating the elderly can generate large private protective returns, while vaccinating younger adults can generate larger public spillover returns through reductions in infections. These disparities motivate the empirical strategy used in this study.

5 Empirical Strategy: Identifying Private and Public 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 the private and public returns to vaccination.

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 how vaccination affects outcomes for the elderly relative to younger adults. This contrast directly measures if vaccination delivers greater private returns in protecting high-risk elderly against severe illness, and/or greater public returns in reducing infections among younger, socially active adults. The TD design also provides a *direct test of the mechanism* by exploiting the differences between the elderly and the younger adults in their propensity for fatal/severe health outcomes and infections. To complement the TD, I also estimate difference-in-differences (DiD) models separately for each age group. These yield the absolute treatment effects of vaccination separately for the elderly and the younger adults.

5.1 Triple-difference design

The first component of my 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. Within this framework, identification comes from differences in vaccination intensity between age groups in the same county and month.

The TD is estimated using a panel regression with a double interaction, Elderly × VaccinationRate, that effectively represents a triple interaction Elderly × (Post-vaccine × High-intensity), 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. 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:

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

where the DiD is as defined as in equation (4), that is explained below. By differencing twice (across time and across counties) to estimate the DiD for each age group, then taking the difference between these DiDs, 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 estimates the differential treatment effect by comparing a more vulnerable group with a less vulnerable one under the same conditions. For private health returns, the elderly serve as the more vulnerable group while for public health returns, the younger adults serve as the more vulnerable group.

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 Triple-difference design tests the mechanism

Beyond serving as an identification strategy, the TD framework also provides direct evidence on the *mechanism* through which vaccines generate heterogeneous effects. By contrasting outcomes for the elderly—who face higher biological vulnerability to severe illness—with those for younger adults—who contribute more to disease transmission—the estimator captures variation along both the *private protective* and *public benefit* channels of vaccine impact. Thus, the TD results not only establish causal heterogeneity but also illuminate the epidemiological and behavioral pathways underlying vaccine effectiveness.

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

While the TD design measure 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. As described above, VaccinationRate captures the effect of Post-vaccine × High-intensity. Therefore, the equation below estimates a DiD:

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

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):

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

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 the county-level 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.

5.4 Reverse Causality from Prior Outcomes?

Figure 4 examines whether reverse causality could plausibly explain the observed variation in vaccine uptake across counties by correlating early vaccination rates (in Mar-2021) with the severity of prior local COVID-19 outcomes (in Dec-2020). Note that even with potential reverse causality, the TD estimates are likely to be unbiased because they compare 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. 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. In fact, statistical tests show that none of the four correlations displayed across the four charts are statistically significant. Thus, figure 4 alleviates concerns about reverse causality biasing the empirical analysis, the DiD estimates in particular.

5.5 Placebo Tests

As additional tests of the validity of the identification strategy, I conduct placebo tests using outcomes unrelated to COVID-19, such as non-COVID deaths; this analysis also utilizes the publicly available data provided by the CDC. Since COVID-19 vaccination should have no causal effect on non-COVID mortality, examining these associations can help us to infer that the main results are not being driven by spurious correlations or omitted factors.

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.

For clarity of exposition, I first focus on identifying the sign and statistical significance of the heterogeneous effects. Discussion of the economic magnitudes and their policy implications is deferred to section 8, after demonstrating the robustness of the estimates to alternative specifications in section 7.

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

I now report in Table 3 the results of the DiD 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 DiD 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 DiD design, I use the methodology developed by Rios-Avila et al. (2022) to implement the robust estimation of DiD in my 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 and denote it as month 0. 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.

We note four findings that emerge from these plots. First, focusing on deaths, we notice that while the estimates for the Average Treatment effect on the Treated (ATT) are strong for the elderly, the effects are relatively muted for the younger adults, which is consistent with the results seen in Tables 2 and 3. 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. All these results are consistent with the 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 interpretation that the DiD estimates are indeed causal, especially because

the reversal in trend post vaccination 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 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 to one year prior. Panel A reports TD estimates, while Panels B and C show separate DiD 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 TD 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, we 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. Moreover, the long-run effects are less likely to be affected by potential reverse causality and are therefore more likely to represent causal estimates. 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.

6.5 Spillovers

A central motivation for this study is to distinguish between the private protective returns to vaccination—reductions in severe health outcomes among those vaccinated—and the public spillover returns that arise when vaccination in one group lowers transmission and indirectly protects others. While the TD and DiD estimates identify the direct effects of vaccination within each age group, they do not by themselves reveal how benefits propagate across groups. Understanding these spillovers is critical for evaluating the full social value of vaccination and for designing optimal allocation strategies. If vaccination of younger adults generates sizable external benefits for the elderly through reduced transmission, the overall welfare ranking of alternative strategies could differ sharply from that implied by private returns alone. To quantify these cross-group effects, I estimate regressions that explicitly relate vaccination rates in one age group to health outcomes in the other, thereby isolating the inter-generational transmission channel.

Table 5 demonstrates how vaccinating one age group (i) creates spillover effects on the health outcomes for the other age group ($-i$). Panel A implements the TD specification:

$$\begin{aligned} y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn_{(-i)c(t-12)} \cdot \Lambda_{65+,i} + \beta_2 Vaccn_{ic(t-12)} \cdot \Lambda_{65+,i} + \\ + \beta_3 Vaccn_{(-i)c(t-12)} + \beta_4 Vaccn_{ic(t-12)} + \beta_5 \Lambda_{65+,i} + \varepsilon_{ct}, \end{aligned} \tag{5}$$

while Panels B and C report the corresponding DiD estimates based on the regression specification:

$$y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn_{(-i)c(t-12)} + \beta_2 Vaccn_{ic(t-12)} + \varepsilon_{ct}. \tag{6}$$

Here i denotes the focal age group (either 65+ or 18–64), and $-i$ the other age group. As all specifications control for the effect of own group vaccination, they estimate the effect of other-group vaccination as a true spillover. Table 5 reports estimates from equations (5) and (6).

Deaths and infections: Across all three panels, vaccination consistently lowers mortality both within and across groups, confirming that vaccines generate substantial protective and spillover benefits. The magnitude of the spillover effects are significantly larger for the elderly than the younger adults. As we seen in Panels B and C, while the spillover effect is 55% greater for the elderly when compared to their own group effect (-0.222 versus -0.143), the spillover effect is about a third of the own group effect for the younger adults (-0.021 versus -0.062). The effects for cases exhibit important similarities and differences to those for deaths. In consistently lowering infection both within and across groups, the results for infections are similar to those for deaths. However, unlike deaths, the magnitude of the spillover effects for infections are significantly larger for the younger adults than the elderly. As we seen in Panels B and C, while the spillover effect is 21% greater for the younger adults when compared to their own group effect (-17.386 versus -14.349), the spillover effect is 11% lower than the own group effect for the elderly (-5.17 versus -5.799). These results reveal a strongly asymmetric spillover pattern: vaccinating the young protects both themselves and the old from fatal outcomes, whereas vaccinating the old primarily protects themselves with relatively smaller spillover effects on the younger adults.

ICU admissions and hospitalizations: The results for ICU and hospitalization outcomes display a striking asymmetry across age groups. Across all three panels, we find that for the elderly, their own vaccination reduces ICU and hospitalizations sharply, while vaccinating younger adults increases these outcomes for the elderly. For younger adults, their own vaccination increases their ICU and hospitalizations, whereas vaccinating the elderly reduces these outcomes. Taken together, these results reveal that vaccinating the elderly consistently lowers severe outcomes across both groups, while vaccinating the younger adults temporarily worsens them. A plausible interpretation is that these effects reflect a combination of differential biological protection and behavioral spillovers. The elderly derive direct clinical benefits from vaccination because of their higher baseline vulnerability to severe disease, producing strong private protective returns. In contrast, when younger adults are vaccinated, however, perceived safety and reduced risk of infection encourage greater mobility, workplace attendance, and inter-generational contact. The elevated exposure may be accounting for the observed rise in ICU and hospitalizations.

Taken together, these findings reveal a consistent pattern across health outcomes. Vaccinating the elderly yields large *private protective returns*, sharply reducing deaths and severe health outcomes for their own group with modest spillovers to younger adults. In contrast, vaccinating younger adults generates substantial *public spillover returns*, lowering infections and mortality among the elderly through reduced transmission, but at the cost of short-term increases in ICU admissions and hospitalizations. This asymmetric pattern highlights the dual role of vaccination as both a private and a public good, and underscores why optimal allocation depends on balancing protection of the vulnerable with interruption of transmission.

6.6 Impact on Severity of Outcomes

A central channel through which vaccination exerts its private health benefits is by mitigating disease severity conditional on infection. To isolate this mechanism, I estimate models using as dependent variables the ratios of deaths to cases and ICU admissions to cases—two severity-adjusted outcomes that capture the conditional probability of progression to severe or fatal illness. These outcomes directly reflect the private protective returns to vaccination, as they measure clinical protection rather than transmission dynamics. Table 6 presents the TD and DiD estimates for these outcomes. The results show that vaccination substantially reduces deaths and ICU admissions per case, especially among the elderly. The TD coefficients in Columns (1) and (4) are negative and highly significant, confirming that vaccine-induced protection against severe disease is stronger for the elderly than for younger adults. The DiD estimates in Columns (2)-(3) and (5)-(6) reveal a similar pattern. Taken together, vaccination lowers the severity of illness conditional on infection with this effect concentrated among those most clinically vulnerable.

7 Robustness Tests And Alternative Explanations

7.1 Placebo Tests Using Non-COVID Deaths

As a falsification exercise, I examine whether COVID-19 vaccination rates are correlated with mortality unrelated to COVID-19. For this purpose, I use state–month level data on non-COVID deaths per million population from CDC as this data is not available at the county level. Using this data, I re-estimate the main specification at different lags of vaccination coverage—contemporaneous, three months, six months, and one year prior. As shown in Table 7, the estimated coefficients are small in magnitude and statistically insignificant across all specifications. These results confirm that the observed effects of vaccination on COVID-19 outcomes are not confounded by unobserved factors correlated with overall mortality, reinforcing the causal interpretation of the main findings.

7.2 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? Table A2 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. I notice in Table A2 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, we focus on the robust evidence that the effect of the two doses of the vaccine cannot be attributed to the effect of the booster dose.

7.3 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 1 and 3. This aggregation helps to internalize most cross-county mobility within metro areas and includes the effects of most proximate spillovers. Table A1 shows the results of this analysis. I find that the results for the TD tests shown in Panel A and the DiD tests for elders and younger adults in Panels B and C respectively remain qualitatively similar. While the coefficients for the DiD estimates for hospitalizations and cases among elders are not statistically significant, their signs remain the same. Similarly, the coefficients for the DiD estimates for ICU admissions and hospitalizations among younger adults are not statistically significant. However, as the coefficients in the TD 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 county 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.

7.4 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? The baseline TD design compares the DiD estimate for the elderly to the corresponding estimate for the younger adults 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 (i) time-invariant county characteristics such as demographics, baseline health infrastructure, baseline exposure risks, political preferences, or persistent differences in health behaviors demographic composition, health infrastructure, and (ii) national-level shocks that vary across months such as the emergence of new variants, 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.

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

First, we consider the possibility of county-specific, time-varying confounders that can bias the TD estimates. 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 for. 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.

Panel A of Table 8 presents results from TD estimates that include fixed effects for each (county, time) pair. Using this rich set of fixed effects, we absorb all county-specific shocks in each month, whether they affect age groups equally or unequally. 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. 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 8. However, the TD coefficient estimate for cases is lower in Panel A of Table 8.

In these tests, the positive coefficient estimate for the average effect of vaccination may seem counterintuitive. 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.

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

A residual concern is that 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, Panel B of Table 8 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.

7.4.3 Triple-difference with state-specific time trends and county and time fixed effects

Panel C of Table 8 reports results from a specification that includes county and time fixed effects together with state-specific linear time trends, which absorb *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. 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. 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.

7.5 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 A3 displays the

results for this robustness analysis. We note that the results remain qualitatively and quantitatively similar to those observed in Tables 2 and 3.

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 5.3, 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. The consistency of the coefficient estimates across the TD and DiD estimates reinforce that the DiD estimates are credibly causal. Therefore, we can infer that the DiD estimates are the most appropriate to use for estimating the economic magnitudes.

8.2 Economic magnitudes

Having established that the DiD estimation provides credible estimates of the causal effect of COVID-19 vaccination, I now assess the economic magnitudes of the heterogeneous effects.

Direct Effects: Using the DiD coefficients estimated in Table 3, I find that 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, 84 ICU admissions , 301 hospitalizations, and 300 cases per million. Compared to the averages for the elderly as seen in Panel A of Table 1, this translates into a decrease of 84% in deaths, 59% in ICU admissions, 23% in hospitalizations, 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, 4.6 ICU admissions, 27 hospitalizations, and 653 cases per million. When compared to the averages for the younger adults, this translates into a decrease of 76% 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.

Spillover Effects: Using Panel A of Table 5, one standard deviation increase in the percentage of the younger adults vaccinated in a county generates a spillover effect of 6 less deaths, 18.9 more ICU admissions, 85.8 more hospitalizations, and 140.1 less cases per million among the elderly, which represent 56% less deaths, 13% more ICU admissions, 6% more hospitalizations, and 6% less cases per million when compared to the average. Thus, even the spillover effects are the most pronounced for the most debilitating outcome—deaths.

A one standard deviation increase in the percentage of the elderly population vaccinated in a county generates a spillover effect of 0.8 less deaths, 20.7 less ICU admissions, 84.1 less hospitalizations, and 652 less cases per million among the younger adults, which represent 32% less deaths, 73% less ICU admissions, 31% less hospitalizations, and 23% less cases per million when compared to their averages. Thus, the spillover effects from both groups are meaningful.

Severity of Outcomes: Compared to the average for deaths per case (= 1.1) and ICU admissions per case (= 29.9) among the elderly, a one standard deviation increase in the percentage of elderly vaccination in a county reduces deaths per case by 136% and ICU admissions per case by 60%. Compared to the average for deaths per case (= 0.25) and ICU admissions per case (= 7.7) among the younger adults, a one standard deviation increase in the percentage of younger adults vaccination in a county reduces deaths per case by 108% and ICU admissions per case by 14%. Thus, vaccination has significant effect in reducing the severity of outcomes conditional on infection as well.

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; only the distribution across age groups is altered. Using the estimated coefficients from the TD and DiD models, health outcomes are simulated one year after roll-out, reflecting the long-run impacts studied above. Crucially, each strategy is evaluated under two modeling assumptions: one that incorporates inter-generational spillovers and another that considers only direct effects within groups. This allows us to isolate the role of externalities in shaping aggregate welfare. Table 10 displays the results of these counterfactual estimations.

9.1 Spillovers Substantially Reshape Predicted Outcomes

Incorporating spillovers dramatically alters the predicted health consequences of vaccine allocation, particularly for younger adults. As shown in Table 10, the direction of effects for younger adults reverses across all three strategies once cross-group externalities are accounted for.

Under the elderly-first strategy (Panel A), ignoring spillovers predicts adverse outcomes for younger adults: deaths increase by 0.46 per million (18.4% when compared to the average), ICU admissions rise by 1.11 per million (3.9%), hospitalizations by 6.73 per million (2.5%) and cases increase by 157.49 per million (5.6%). This is expected because lower vaccination of the younger adults leads to a direct negative impact on health outcomes as seen in Tables 2 and 3. However, when spillovers are included, the picture *flips*: younger adults experience broad improvements. Most strikingly, deaths fall by 0.15 per million (6%), ICU admissions fall by 15.46 per million (54.8%), hospitalizations drop by 58.98 per million (21.8%), and cases decline by 363.61 per million (13%). This reversal highlights the powerful positive externalities on younger adults generated by vaccinating the elderly first.

Conversely, under the younger-adult-first strategy (Panel B), excluding spillovers suggests modest gains for younger adults—for example, a 28% reduction in deaths. But including spillovers reveals large negative externalities: because the elderly receive fewer doses relative to their clinical risk, mortality and severe outcomes surge among them. Once the negative spillovers from lower vaccination of the elderly are accounted for, all outcomes for younger adults worsen: deaths rise by 0.22 per million (+8.8%), ICU admissions increase by 23.58 per million (+83.6%), hospitalizations rise by 89.91 per million (+33.2%), and cases jump by 553.81 per million (+19.8%).

Similarly, under proportional allocation (Panel C), the apparent benefits for younger adults vanish when spillovers are included. While this strategy spreads doses evenly, it fails to achieve

sufficient coverage among the high-risk elderly—producing avoidable deaths that generate downstream costs for the entire system. When spillovers are incorporated, younger adults see increases in deaths (+4%), ICU admissions (+35.5%), hospitalizations (14.1%) and cases (+8.5%).

Even when we focus on the health outcomes of the elderly, spillovers from vaccinating the younger adults are significant. Incorporating the spillovers into the counter-factual estimations reduces the impact on deaths per million and infections per million by 63% and 44% respectively while increasing ICU admissions and hospitalizations by 21% and 27% respectively.

9.2 Elderly-First Remains Optimal Despite Strong Spillovers

Despite the powerful public-good benefits of vaccinating younger adults, particularly in reducing transmission, the *elderly-first strategy minimizes overall mortality and critical-care burden.* This result holds even after fully accounting for spillovers. Two forces explain this outcome:

1. *Clinical vulnerability*: Each vaccine dose yields a much larger mortality reduction among the elderly, who face disproportionately higher baseline risks of severe illness and death.

2. *Demographic leverage*: Because the elderly constitute a smaller share of the population (17.9% on average across counties), high vaccine coverage for the elderly can be achieved quickly with limited supply. In fact, a 1% reallocation of vaccine doses from the younger adults to the elderly increases elderly coverage by almost 5%, while the reverse shift provides only marginal protection to younger cohorts. Consequently, even though the absolute spillover effects from younger to older adults are larger, vaccinating the elderly first produces greater overall welfare gains when supply is limited.

Thus, the elderly-first strategy exploits differences in both clinical vulnerability and demographic efficiency.

9.3 Economic Magnitudes and Policy Implications

The simulated effects are sizable. The magnitude of the simulated effects is substantial. Under elderly-first allocation with spillovers, elderly deaths fall by 21.6% (2.31 per million), ICU admissions drop by 50.1% (71.19 per million), hospitalizations decline by 20.1% (267.55 per million), and infections drop by 4.7% (118.77 per million). As described above, younger adults also benefit. These gains come at no cost to broader transmission control.

Taken together, these simulations show that optimal vaccine policy depends jointly on the magnitudes of private protective returns, public spillover benefits, and population asymmetries. Vaccinating younger adults generates larger externalities, but vaccinating the elderly saves more lives per dose because of their higher vulnerability and smaller cohort size. By embedding empirically estimated spillovers into the simulations, this analysis moves beyond epidemiological heuristics and provides an integrated framework for designing age-targeted vaccination policies in future pandemics.

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—and when these effects spill across population groups. While the empirical analysis exploits detailed U.S. county-level data, the underlying tension between protecting the most vulnerable and minimizing transmission is conceptually general. The results show that vaccination generates both strong *private protective returns*, concentrated among the elderly, and sizable *public spillover benefits*, particularly from vaccinating younger adults. Crucially, these inter-generational spillovers are not peripheral—they fundamentally alter the welfare calculus. Yet, even after accounting for them, prioritizing the elderly remains the optimal policy in the U.S. context. The explanation lies in demographic leverage: vaccinating a smaller, high-risk group quickly produces disproportionately large reductions in disease burden.

10.1 Application to Other Country Settings

The U.S. context provides a useful benchmark, but cross-country demographic and institutional differences can shift the balance between private protection and public spillovers. The U.S. has an older population than many low- and middle-income countries but younger than most European or East Asian societies. In younger societies, the transmission externalities from vaccinating working-age cohorts are likely stronger, potentially raising the social return to vaccinating younger adults first. Conversely, in aging populations with high baseline mortality risk, the private protective returns to vaccinating the elderly remain large—but the demographic leverage that makes elderly-first optimal in the U.S. weakens substantially.

In the U.S., reallocating 1% of available doses from younger adults to the elderly increases elderly vaccine coverage by nearly 5%, producing large welfare gains. In contrast, in countries such as Japan—where the elderly comprise almost 30% of the population—the same 1% reallocation raises elderly coverage by only slightly more than 2%. This sharp decline in demographic leverage fundamentally alters the welfare calculus. When the elderly represent such a large share of the population, the efficiency advantage of concentrating doses on them diminishes, while the relative importance of transmission spillovers from vaccinating younger cohorts grows. Thus, in very elderly societies, vaccinating the elderly first may no longer dominate other strategies once spillovers are fully incorporated, as the aggregate benefits of slowing transmission among the larger working-age population could outweigh the smaller incremental protection achieved for the elderly.

10.2 Applicability beyond COVID-19

Although this study focuses on COVID-19 and spans multiple variants (Alpha, Delta, and Omicron), its insights extend beyond this specific pathogen. The triple-difference design isolates the average treatment effects of vaccination under real-world heterogeneity in variant prevalence, vaccine efficacy, and behavioral response. Future pandemics may differ in both age-specific severity and transmission dynamics. For instance, a pathogen that disproportionately affects younger adults

would invert the prioritization calculus observed here, making public spillovers dominate private protection. What remains general, however, is the analytical framework: policymakers must measure and compare private protective returns and public spillover benefits in real time, recalibrating vaccine allocation as epidemiological evidence evolves. The approach of integrating causal estimates of spillovers into counterfactual policy simulations provides a framework for such dynamic re-optimization.

10.3 Macroeconomic dimensions of the vaccination trade-off

The heterogeneous effects documented here also carry macroeconomic significance. Vaccination spillovers that reduce infections among younger adults, who are engaged in the labor force, can yield measurable gains in labor supply, productivity, and economic stability. Although this paper does not directly estimate those channels, prior research establishes strong links between population health and macroeconomic output (Perkins et al., 2023; Deb et al., 2021). By mitigating illness-induced absenteeism, sustaining consumer activity, and reducing healthcare utilization, widespread vaccination produces positive economic externalities that complement its epidemiological effects. At the same time, protecting the elderly alleviates pressure on intensive-care systems and long-term care expenditure, reinforcing fiscal stability. Hence, vaccine allocation embodies a dual optimization problem: maximizing health welfare while preserving economic resilience.

In sum, this study underscores that effective vaccination policy must integrate three dimensions: (i) heterogeneous biological returns across demographic groups, (ii) public spillover effects arising from transmission externalities, and (iii) population structure and capacity constraints that mediate both.

11 Conclusion

This paper provides new empirical evidence on the heterogeneous and interdependent benefits of COVID-19 vaccination across age groups, highlighting a central policy dilemma between private, protective returns and public benefits. Using comprehensive U.S. county-level panel data and complementary quasi-experimental designs, I show that vaccination yields long-run health benefits that differ sharply by age group and health outcome. Vaccinating the elderly delivers large private protective gains while vaccinating younger adults generates substantial public benefits by curbing infections across groups. The spillover analysis further demonstrates that vaccinating younger adults significantly improves outcomes for the elderly through reduced transmission, whereas vaccinating the elderly produces smaller but still meaningful indirect benefits for younger adults.

Counterfactual simulations reveal that accounting for these spillovers reshapes predicted outcomes but does not overturn the optimality of prioritizing the elderly in the U.S. context. The advantage of the elderly-first strategy arises from demographic leverage: because the elderly represent a smaller share of the population, reallocating doses toward them rapidly achieves high coverage and large welfare gains. Optimal vaccination policy must therefore balance biological

vulnerability, demographic structure, and social externalities—a framework that extends beyond COVID-19 to 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.

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.

Hendren, N. and B. Sprung-Keyser (2020). A unified welfare analysis of government policies. *The Quarterly Journal of economics 135*(3), 1209–1318.

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.

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 2-doses of vaccine, (2) population share receiving at least 1-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 1-dose, 2-doses and the booster dose. This data is provided for the following two age groups: 18+ years and 65+ years.

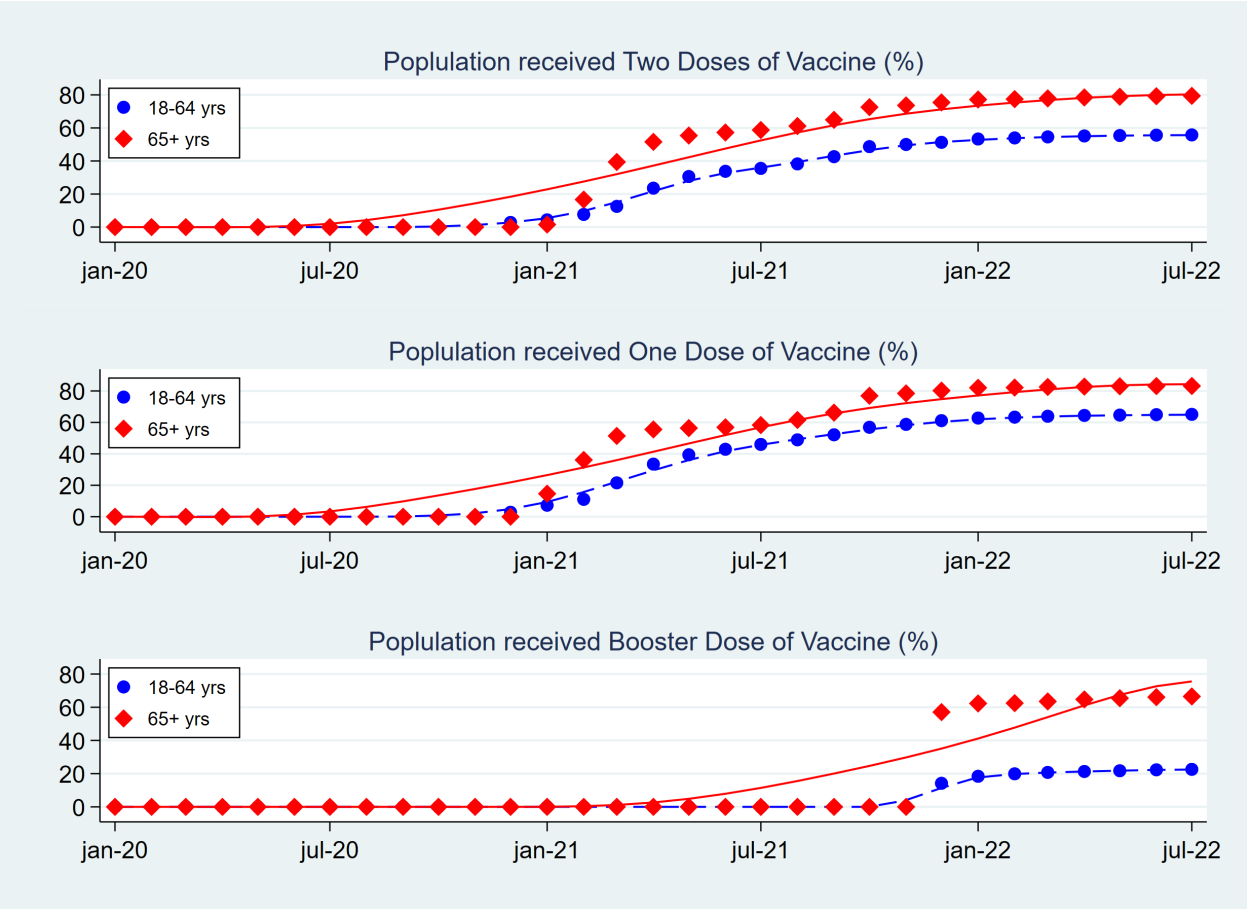


Figure 2: Geographic and Temporal Variation in Two-Dose COVID-19 Vaccine Coverage by Age Group Across U.S. Counties from Jan 2020 to Jul 2022

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 1-dose, 2-doses and the booster dose. This data is provided for the following two age groups: 18+ years and 65+ years.

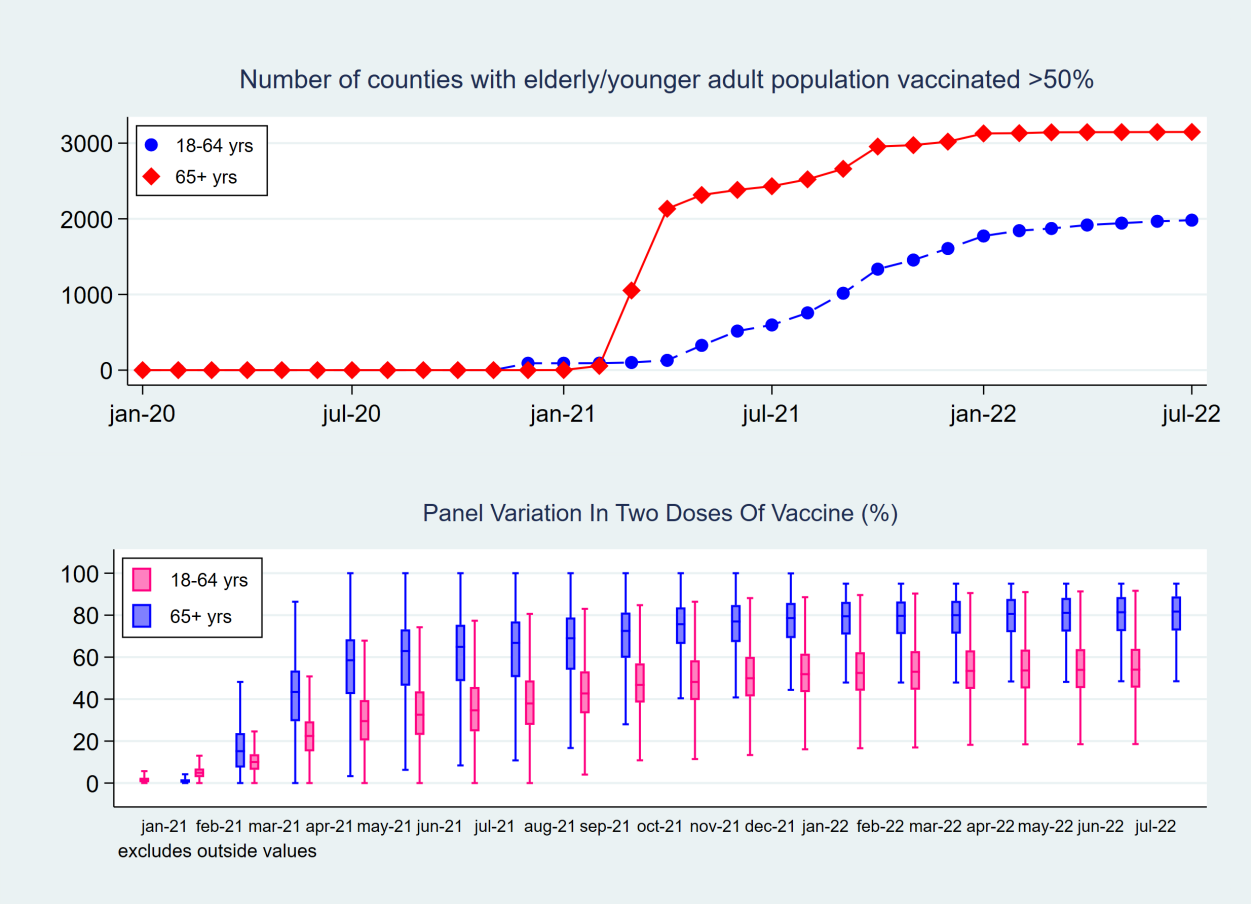


Figure 3: Trends in COVID-19 Health Outcomes by Age Group (U.S. Counties, Jan 2020–Jul 2022)

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 (18–64 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.

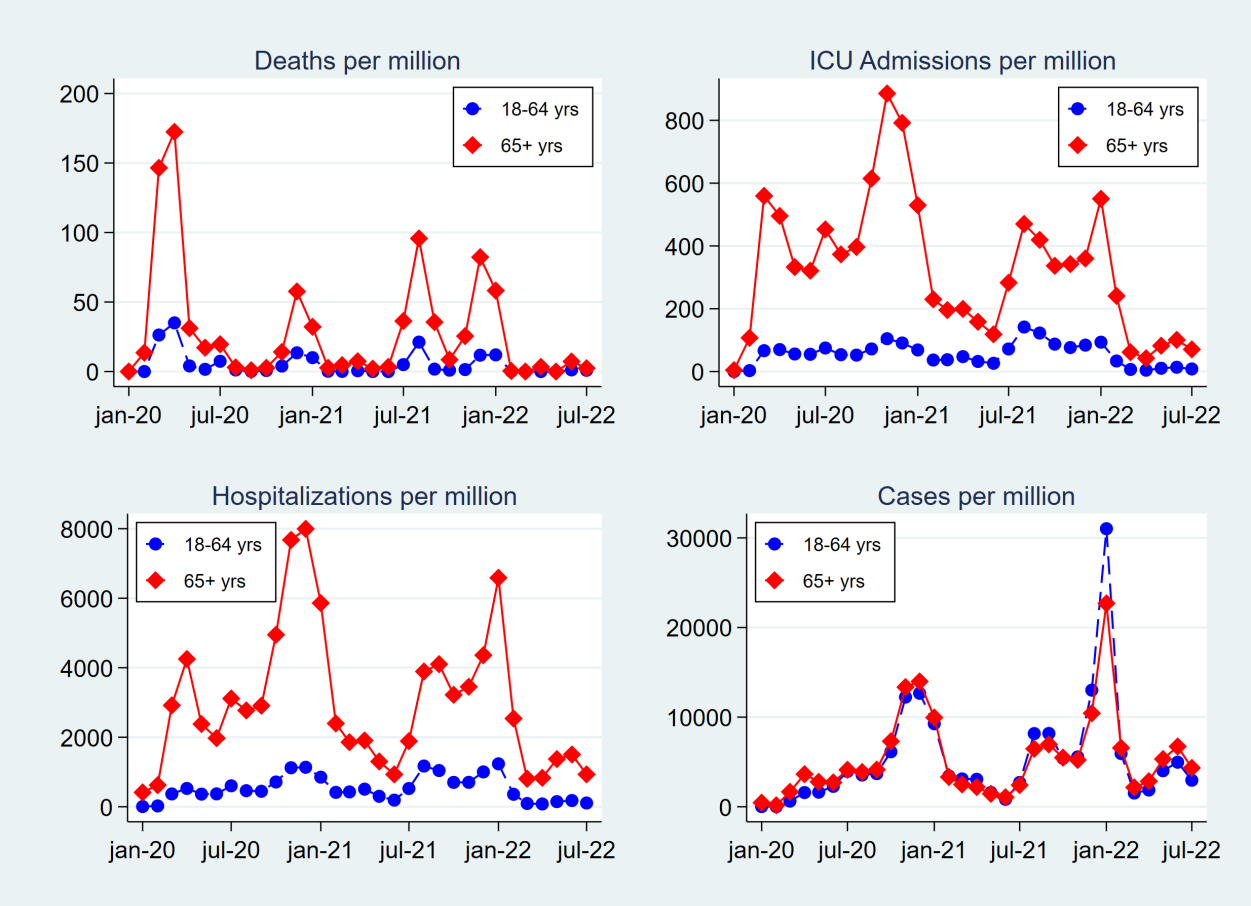


Figure 4: Pre-Vaccine COVID-19 Health Outcomes and Mar-2021 Vaccination Rates: Assessing Reverse Causality

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 1-dose (blue circles) and 2-doses (red diamonds). The top panels shows the (lack of) correlation of vaccination—both 1-dose and 2-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 1-dose and 2-doses—for the younger adults (18–64 years) in Mar-2021 with cases per million three months earlier in Dec-2020.

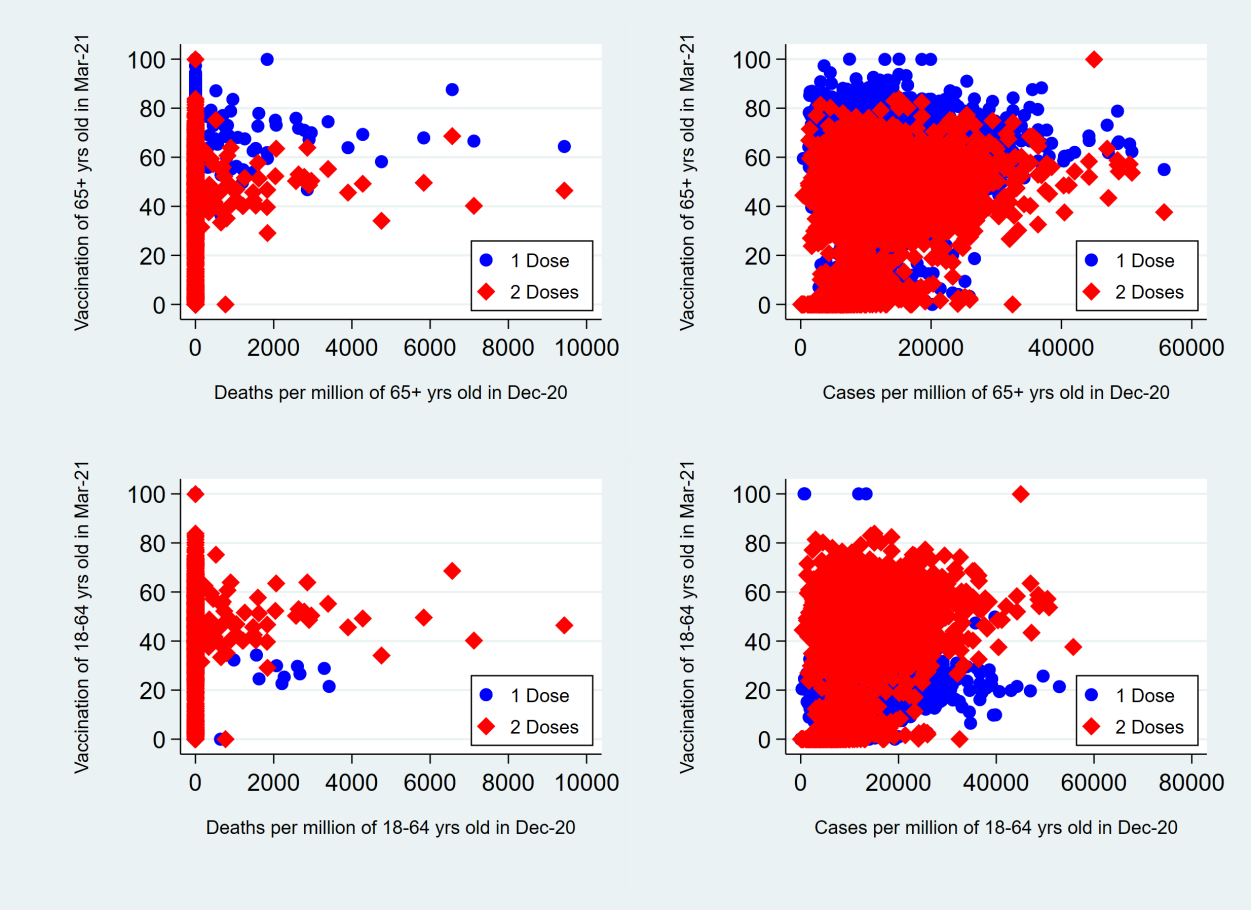
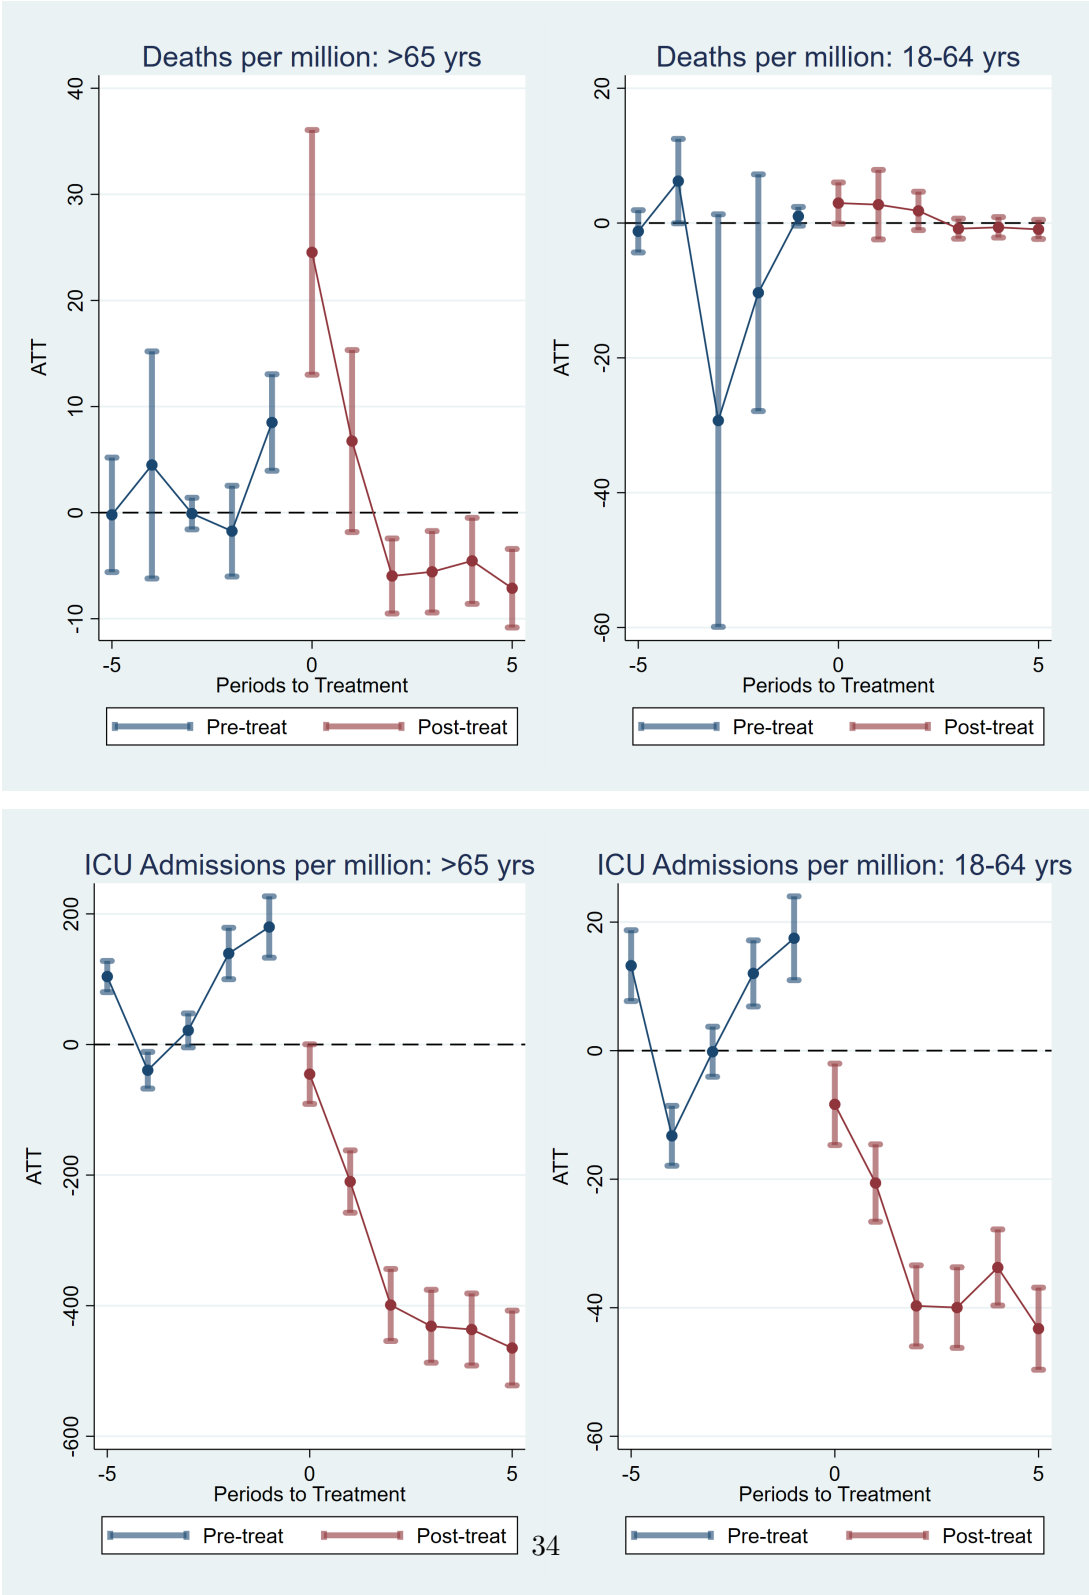
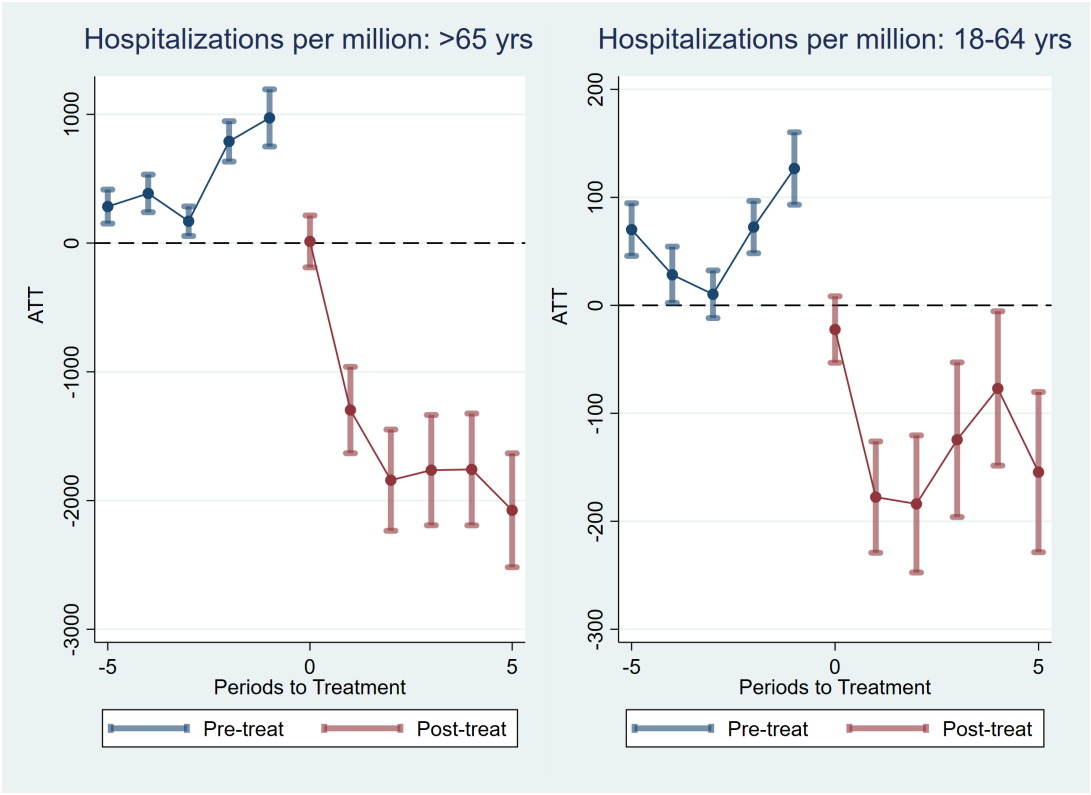


Figure 5: Dynamic effects estimated using the methodology in Rios-Avila et al. (2022) of "not yet treated" counties as control group

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.




Hospitalizations per million: >65 yrs
ATT
1000
0
-1000
-2000
-3000
-5
0
5
Periods to Treatment
Pre-treat
Post-treat
Hospitalizations per million: 18-64 yrs
ATT
200
100
0
-100
-200
-300
-5
0
5
Periods to Treatment
Pre-treat
Post-treat

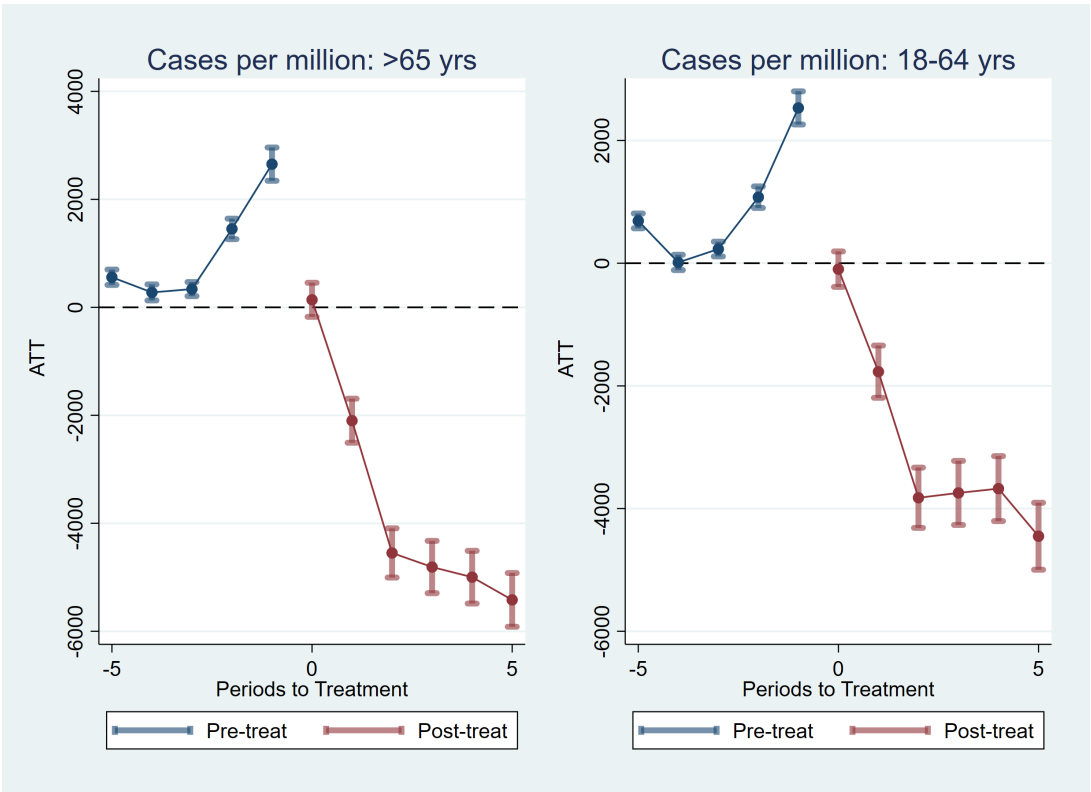

Cases per million: >65 yrs
ATT
4000
2000
0
-2000
-4000
-6000
-5
0
5
Periods to Treatment
Pre-treat
Post-treat
Cases per million: 18-64 yrs
ATT
2000
0
-2000
-4000
-6000
-5
0
5
Periods to Treatment
Pre-treat
Post-treat

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 2-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 2-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 2-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. *** denotes statistical significance at the 99% level.

	(1)	(2)	(3)	(4)
Dependent variable per million:	Deaths	ICU	Hospital	Cases
Population received 2-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 2-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 and Time FE	Yes	Yes	Yes	Yes

Table 3: Effect of 2-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 2-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. *** denotes statistical significance at the 99% level.

	(1)	(2)	(3)	(4)
Dependent variable per million:	Deaths	ICU	Hospital	Cases
Panel A: Elderly (65+ yrs)				
Population received 2-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 and Time FE	Yes	Yes	Yes	Yes
Panel B: Younger Adults (18–64 yrs)				
Population received 2-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 and 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% and 95% 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 and 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 and 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 and Time FE	Yes	Yes	Yes	Yes

Table 5: Spillover Effects of Vaccinating Younger Adults on Elderly and Vice-versa

This table displays the results of panel regressions examining the impact of COVID-19 vaccination of elderly (65+ yrs) on health outcomes of the younger adults (18–64 yrs) and vice-versa. Panel A implements the following OLS regression for the triple-difference estimate:

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

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

$y_{ict} = \beta_c + \beta_t + \beta_1 Vaccn_{(-i)c(t-12)} + \beta_2 Vaccn_{ic(t-12)} + \varepsilon_{ct}$

where $-i$ denotes the age group that is not age group i. Monthly data from January 2020 to Jul-2022 are obtained from the CDC. ***, **, and * denote statistical significance at the 99%, 95% and 90% levels respectively using robust standard errors clustered at the county-level in parentheses.

Dependent variable per million:	(1) Deaths	(2) ICU	(3) Hospital	(4) Cases
Panel A: Triple-difference estimates				
Other age group recd. 2-doses 1-year before (%)	-0.142***	1.170***	2.259***	7.943***
×Dummy for 65+ years	(0.043)	(0.119)	(0.764)	(2.756)
Own age group recd. 2-doses 1-year before (%)	-0.022	-2.757***	-12.766***	4.276
×Dummy for 65+ years	(0.039)	(0.158)	(0.806)	(2.909)
Other age group recd. 2-doses 1-year before (%)	-0.013	-0.972***	-1.660***	-18.369***
	(0.023)	(0.077)	(0.552)	(1.638)
Own age group recd. 2-doses 1-year before (%)	-0.129***	0.645***	2.799***	-9.093***
	(0.032)	(0.073)	(0.526)	(2.108)
Dummy for 65+ years	9.569***	126.242***	1,142.840***	-374.073***
	(0.880)	(5.902)	(26.928)	(15.024)
Observations	199,950	199,950	199,950	199,950
Adjusted R-squared	0.064	0.217	0.309	0.431
Panel B: Difference-in-difference estimates for the Elderly (65+yrs)				
Other age group recd. 2-doses 1-year before (%)	-0.222***	0.698***	3.166***	-5.170***
	(0.049)	(0.120)	(0.679)	(1.946)
Own age group recd. 2-doses 1-year before (%)	-0.143***	-2.533***	-9.385***	-5.799***
	(0.041)	(0.153)	(0.880)	(1.522)
Observations	99,975	99,975	99,975	99,975
Adjusted R-squared	0.072	0.297	0.394	0.420
Panel C: Difference-in-difference estimates for the Younger Adults (18–64 yrs)				
Other age group recd. 2-doses 1-year before (%)	-0.021*	-0.552***	-2.242***	-17.386***
	(0.012)	(0.034)	(0.200)	(1.705)
Own age group recd. 2-doses 1-year before (%)	-0.062***	0.144***	0.232	-14.349***
	(0.018)	(0.023)	(0.183)	(2.022)
Observations	99,975	99,975	99,975	99,975
Adjusted R-squared	0.060	0.313	0.396	0.439
County FE and Time FE	Yes	Yes	Yes	Yes

Table 6: Effects of Vaccination on Disease Severity

This table displays the results of panel regressions examining the impact of COVID-19 vaccination on severity of health outcomes conditional on infection. Columns 1 and 4 implement the following OLS regression for the triple-difference estimate:

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

Columns 2-3 and 5-6 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. ***, **, and * denote statistical significance at the 99%, 95% and 90% levels respectively using robust standard errors clustered at the county-level in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Deaths / Cases			ICU / Cases		
	TD	DiD (65+)	DiD (18–64)	TD	DiD (65+)	DiD (18–64)
Popln. recd. 2-doses 1-yr before (%) × Dummy for 65+ years	-0.02***			-0.32***		
	(0.00)			(0.03)		
Popln. recd. 2-doses 1-yr before (%)	0.00***	-0.04**	-0.01***	0.15***	-0.48***	-0.04***
	(0.00)	(0.02)	(0.00)	(0.02)	(0.03)	(0.01)
Dummy for 65+ years	1.03***			24.47***		
	(0.16)			(1.55)		
Observations	199,950	99,975	99,975	199,950	99,975	99,975
Adjusted R-squared	0.29	0.07	0.05	0.27	0.27	0.18
County FE and Time FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 7: Placebo Tests - Effects of Vaccination on Non-COVID Deaths

This table reports placebo estimates testing whether COVID-19 vaccination rates predict mortality unrelated to COVID-19. The dependent variable is the number of non-COVID deaths per million population at the state–month level from January 2020 to July 2022 in the following regression:

$y_{st} = \beta_s + \beta_t + \beta_1 Vaccn_{s(t-k)} + \varepsilon_{st}, k = 0, 3, 6, 12.$

Standard errors shown in parentheses are clustered at the state level.

	(1)	(2)	(3)	(4)
Dependent Variable:	Non-COVID Deaths per million			
Two-dose vaccination in the same month	-23.84			
	(72.33)			
Two-dose vaccination three months before		-5.64		
		(83.53)		
Two-dose vaccination six months before			44.78	
			(76.64)	
Two-dose vaccination one year before				39.18
				(87.33)
Observations	1,550	1,550	1,550	1,550
Adjusted R-squared	0.84	0.84	0.84	0.84
State and Time FE	Yes	Yes	Yes	Yes

Table 8: 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.

Dependent variable per million:	(1) Deaths	(2) ICU	(3) Hospital	(4) Cases
Panel A: Triple-difference using FE for each (County, Time) pair				
Population received 2-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 2-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 2-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 2-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 2-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 2-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 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 8	-0.16***	-1.59***	-10.51***	12.22***
(County*Age Group) FE and Time FE from Panel B of Table 8	-0.12***	-1.75***	-11.35***	23.04***
State-specific time trends + county FE and time FE from Panel C of Table 8	-0.16***	-1.40***	-9.65***	19.51***
Controlling for booster dose from Table A2	-0.38***	-2.05***	-22.05***	68.98***
MSA level tests from Table A1	-0.57***	-2.87***	-20.29***	16.46***
Excluding observations in year 2020 from Table A3	-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***

Table 10: Counterfactual Simulations: Health Outcomes Under Alternative Vaccination Strategies, With and Without Spillovers

This table presents simulated average changes in COVID-19 health outcomes per million population under three vaccination strategies—prioritizing the elderly in Panel A, prioritizing younger adults in Panel B, and proportional allocation in Panel C—under two scenarios: with and without accounting for inter-group spillovers. All the results are reported for the elderly (65+ years) and for the younger adults (18–64 years). All simulations use the estimated difference-in-difference and spillover coefficients, setting statistically insignificant effects to zero. Negative values represent health improvements when compared to the actual health outcomes. All the values reported are in health outcomes per million; the numbers in parentheses denote the percentage change when compared to the respective average.

Panel A: Prioritize Elderly (65+ yrs)

Age Group	Spillovers?	Deaths per million	ICU per million	Hospital per million	Cases per million
Elderly	With	-2.31 (-21.6%)	-71.19 (-50.1%)	-267.55 (-20.1%)	-118.77 (-4.7%)
Elderly	Without	-6.31 (-59%)	-58.92 (-41.5%)	-211.49 (-15.9%)	-210.44 (-8.3%)
Younger Adults	With	-0.15 (-6%)	-15.46 (-54.8%)	-58.98 (-21.8%)	-363.61 (-13%)
Younger Adults	Without	0.46 (18.4%)	1.11 (3.9%)	6.73 (2.5%)	157.49 (5.6%)

Panel B: Prioritize Younger Adults (18–64 yrs)

Age Group	Spillovers?	Deaths per million	ICU per million	Hospital per million	Cases per million
Elderly	With	3.52 (32.9%)	108.55 (76.4%)	408 (30.7%)	180.88 (7.1%)
Elderly	Without	9.62 (89.9%)	89.83 (63.3%)	322.42 (24.3%)	320.82 (12.7%)
Younger Adults	With	0.22 (8.8%)	23.58 (83.6%)	89.91 (33.2%)	553.81 (19.8%)
Younger Adults	Without	-0.7 (-28%)	-1.7 (-6%)	-10.29 (-3.8%)	-240.96 (-8.6%)

Panel C: Proportional Allocation

Age Group	Spillovers?	Deaths per million	ICU per million	Hospital per million	Cases per million
Elderly	With	1.51 (14.1%)	46.04 (32.4%)	173 (13%)	77.13 (3%)
Elderly	Without	4.09 (38.2%)	38.14 (26.9%)	136.88 (10.3%)	136.2 (5.4%)
Younger Adults	With	0.1 (4%)	10 (35.5%)	38.17 (14.1%)	236.06 (8.5%)
Younger Adults	Without	-0.29 (-11.6%)	-0.71 (-2.5%)	-4.3 (-1.6%)	-100.7 (-3.6%)

Online Appendix For
Designing Pandemic Vaccination:
Private and Public Returns in Optimal Allocation

Table A1: 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 2-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 2-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 (65+yrs)				
Population received 2-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 (18–64 yrs)				
Population received 2-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 A2: Incremental Effect Of 2-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. *** denotes statistical significance at the 99% level.

Dependent variable per million:	(1) Deaths	(2) ICU	(3) Hospital	(4) Cases
Population received 2-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 2-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 and Time FE	Yes	Yes	Yes	Yes

Table A3: 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% and 95% levels respectively.

	(1)	(2)	(3)	(4)
Dependent variable per million:	Deaths	ICU	Hospital	Cases
Panel A: Triple-difference estimates				
Population received 2-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 2-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 and Time FE	Yes	Yes	Yes	Yes
Panel B: Difference-in-difference estimates for the Elderly (65+yrs)				
Population received 2-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 and Time FE	Yes	Yes	Yes	Yes
Panel C: Difference-in-difference estimates for the Younger Adults (18–64 yrs)				
Population received 2-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 and Time FE	Yes	Yes	Yes	Yes