

NBER WORKING PAPER SERIES

DIRECT AND INDIRECT EFFECTS OF VACCINES:
EVIDENCE FROM COVID-19

Seth M. Freedman
Daniel W. Sacks
Kosali I. Simon
Coady Wing

Working Paper 30550
<http://www.nber.org/papers/w30550>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2022, revised August 2023

We thank Dr. Sarah Caddy and Dr. Buddy Creech for helpful background on immunology. We are grateful to our discussants Onur Altindag, Kitt Carpenter, and Emily Lawler, Ajin Lee, as well as, Timo Hener, Scott Imberman, Frederik Lyngse, and Emily Oster, and audiences at Aarhus University, Allied Social Science Association Annual Meeting, ASHEcon Annual Conference, Copenhagen Business School, DePaul University, Indiana University, Junior Health Economics Summit, Korean Women Economists Network (KWEN), NBER Summer Institute, University of Wisconsin - Madison, University of Wisconsin - Milwaukee, and Washington University in St. Louis for helpful comments. Mallory Dryer, Aneka Fish, Dario Salcedo Monroy, Maddie Mustaine, Anna Seng, and Maddie Yowziak provided excellent research assistance. We acknowledge the Indiana University Pervasive Technology Institute for providing HPC resources that have contributed to the research results reported in this paper (<https://pti.iu.edu>), and we thank Haile Shone for helpful data management. This project was partially funded by the Vice Provost for Research through the Research Equipment Fund. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Seth M. Freedman, Daniel W. Sacks, Kosali I. Simon, and Coady Wing. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Direct and indirect effects of vaccines: Evidence from COVID-19
Seth M. Freedman, Daniel W. Sacks, Kosali I. Simon, and Coady Wing
NBER Working Paper No. 30550
October 2022, revised August 2023
JEL No. I0,I28

ABSTRACT

We estimate direct and indirect vaccine effectiveness, and assess how far the infection-reducing externality extends from the vaccinated, a key input to policy decisions. Our empirical strategy uses nearly universal microdata from a single state and relies on the six-month delay between 12- and 11 year-old COVID vaccine eligibility. Vaccination reduces cases by 80 percent, the direct effect. This protection spills over to close contacts, producing a household-level indirect effect about half as large as the direct effect. However, indirect effects do not extend to schoolmates. Our results highlight vaccine reach as an important aspect of policy towards infectious disease.

Seth M. Freedman
School of Public and Environmental Affairs
1315 E. 10th St.
Bloomington, IN 47405
US
freedmas@indiana.edu

Daniel W. Sacks
975 University Ave
Madison, WI 53715
dan.sacks@wisc.edu

Kosali I. Simon
O'Neill School of Public and Environmental Affairs
Indiana University
1315 East Tenth Street
Bloomington, IN 47405-1701
and NBER
simonkos@indiana.edu

Coady Wing
Indiana University
1315 E 10th St
Bloomington, IN 47405
cwing@indiana.edu

I Introduction

Effective vaccines reduce illness among people who are vaccinated, a *direct effect* that mitigates the damages from pandemic and endemic diseases. Some vaccines also have *indirect effects* by reducing the transmission of disease from the vaccinated person to others, either by reducing primary infection or by reducing contagiousness given infection. These indirect effects may be critical to the ability of vaccines to end pandemics and eradicate diseases. When vaccines produce indirect effects, the benefits of vaccination do not only accrue to the vaccinated, making vaccines a textbook example of a good with positive externalities (Gruber, 2005). Vaccine externalities provide a neoclassical justification for common but controversial government interventions in vaccine markets, for COVID and other viral diseases. These government interventions in vaccine markets include billions of dollars for research, development, and deployment, as well as heavier-handed measures such as vaccine mandates or passports, which require vaccination for continued employment, travel, or entry to social venues.¹ Non-government organizations such as employers and higher education institutions have also implemented vaccine requirements for workers and customers. Some national and state governments have adopted vaccine passports or mandates, while others have moved in the other direction by banning vaccine mandates or making it difficult or illegal for private parties to compel vaccination as a condition of employment or participation in various events or services.

The case for vaccine market interventions depends in part on the magnitude and *reach* of the vaccine effects. If the reach is narrow—one person’s vaccination primarily benefits her and her closest contacts—then private contracting may be sufficient to achieve near-optimal levels of vaccination. Narrow reach implies that while vaccine mandates may be unnecessary, bans on private mandates are harmful. If the reach is far in the sense that vaccination protects against brief contacts with coworkers from other officers, bus travelers, or restaurant diners, then contracting frictions likely impede optimal vaccination. In that case, policies like mandates, subsidies, or other inducements are apt to be welfare improving.

Vaccine effects—direct or indirect—are difficult to measure. Randomized clinical trials provide in-

¹Government spending supporting development and subsidized doses of the COVID vaccine came to more than \$30 billion as of 2023 (Kates et al., 2023).

ternally valid estimates of direct effects, but these estimates may not generalize outside the clinical trial. Clinical trial populations are not necessarily representative of the population, in terms of demographics (e.g. Hall (1999)), expected benefits (Chan and Hamilton, 2006; Malani, 2008), or study site (Allcott, 2015). Blinded clinical trials do not capture behavioral responses such as increased risk taking that may offset the health benefits of a vaccine in the field (Chan et al., 2016). Finally, viral variants that evade the vaccine immune response may reduce effectiveness. Furthermore, vaccine trials are not designed to detect indirect effects. Direct and indirect effects can in principle be estimated in observational settings; however, without random assignment, vaccination will typically be confounded by unmeasured health and risk characteristics. Identifying indirect effects further requires identifying peers and isolating exogenous variation in their vaccination status, conditional on own vaccination.

In this paper we overcome these challenges to identifying direct vs indirect effects using unique, detailed microdata on vaccination and health from Indiana to provide quasi-experimental evidence on the direct and indirect effects of the COVID vaccine, as well as an assessment of reach. Our core identification strategy uses age-based variation in eligibility, along with variation in peer ages, to separately identify direct and indirect effects of the vaccine in two important settings: households and schools. We find large direct effects, with vaccine effectiveness of about 80%. Indirect effects are context dependent. Within a household, indirect effects of the vaccine are substantial: one household member's vaccination protects other household members by about half as much as the direct effect. In contrast, we find small and insignificant indirect effects of the vaccine in school settings. Together, these results suggest indirect effects are substantial, though have limited reach, and provide valuable frameworks for future public health policy towards different infectious diseases.

Our empirical strategy takes advantage of the age-based roll out of the COVID vaccines. Children aged 12-16 became vaccine-eligible in May of 2021, while children aged 5-11 were ineligible until November. Thus 12 year-olds were eligible six months sooner than 11 year-olds. We study the effect of earlier eligibility on the 12 year-olds themselves, as well as on their household members and their schoolmates. We use data comprising the near-universe of vaccinations and lab-based COVID tests in Indiana, combined with

health care records and demographic information for a large subset of the state population.

Our first contribution is to estimate the direct effects of vaccination on COVID incidence among the vaccinated. Using a difference-in-differences design, we show that early eligibility increased vaccination rates among 12 year-olds by 22 percentage points, and reduced monthly COVID incidence among 12 year-olds by 0.3 percentage points. In an instrumental variables framework, our results imply that vaccination reduces monthly COVID incidence by about 1.3 percentage points, 80 percent of the control complier mean. We detect no adverse health effects of vaccination. This result confirms evidence from clinical trials that the COVID vaccines are highly effective, but extends that evidence to a field setting where the Delta variant was prominent and people knew their own vaccination status.

Our primary contribution is to estimate indirect effects of the vaccines, measuring how much and how far vaccine protection extends beyond the vaccinated. To identify household-level indirect effects, we compare adults residing with 12 year-olds to adults residing with 11 year-olds, in the period after 12 year-old vaccine availability, relative to the period before. Because these adults are all vaccine-eligible, similar on baseline characteristics, and exhibit identical pre-treatment trends, this comparison isolates the effect of one household member's vaccine eligibility on other household member's COVID incidence.

We find substantial indirect effects within households. Our results imply that, for each COVID case prevented by vaccinating a 12 year-old, another half a case is prevented among her household members. This indirect effect reflects the combined effect of two separate vaccine effect channels: reduction in own infection and reduction in transmission conditional on infection, what VanderWeele et al. (2012) call "contagion" and "infectiousness" effects. Alternative approaches focus on the second channel by measuring how infection reduces the "secondary attack rate" (Halloran et al., 2003). This requires conditioning on infection, raising endogeneity issues our approach avoids.

This household-level indirect effect is specific to our context and to our instrument-complier population. As the compliers are adults who vaccinate their children upon eligibility, the complier vaccination rate is high, 80 percent, versus 35 percent for never takers. Thus our design measures spillovers to an already vaccinated complier population. We speculate that spillovers could be larger in populations with

lower take-up. We further caution that we measure households only imperfectly, although sensitivity analyses suggest any bias from the resulting measurement error is likely small. Both concerns imply that we may underestimate household indirect effects in the population.

While we find large household-level indirect effects, we find near-zero effects in schools. Our identification strategy takes advantage of the interaction between school assignment rules and vaccine eligibility. At the start of the 2021 school year, all seventh grade and older students are vaccine eligible, but all fifth grade and younger students are vaccine ineligible; sixth graders are mostly ineligible. But some sixth graders attend middle schools with vaccine eligible older students. Other sixth graders attend elementary schools with vaccine ineligible younger students. To estimate the effect of going to school with more vaccinated peers, we compare sixth graders in middle schools to sixth graders in elementary schools in a difference-in-differences framework. Middle and elementary school sixth graders exhibited near-identical trends in COVID incidence before the vaccine roll out. After roll out, middle school sixth graders experience a 23 percentage point differential increase in the share of their schoolmates who are vaccinated. However, this increase in school-wide vaccination led to a small, positive (i.e. wrong signed), and statistically insignificant effect on their COVID incidence.

This result suggests attending school with more vaccine-eligible children does not provide protective benefits. One possible explanation is that different grades mix too infrequently for seventh grade vaccinations to protect sixth graders, especially during times of social distancing in non-household settings. We rule out this explanation with our final identification strategy. We use a regression discontinuity design to compare students born just before and after the cutoff to enter seventh grade in the Fall of 2021. The cut-off children differ sharply in the share of their grade mates who are vaccine eligible. We find no protective effect of this increase in COVID-vaccine eligible grade mates, although low take-up and limited sample sizes at the discontinuity limit our ability to precisely estimate discontinuities.

Overall, we find heterogeneous indirect effects: clear and substantial at the household level, small and insignificant at the school level. We argue that a key explanation is that there is greater mixing among household members than among schoolmates. We rule out the other leading explanation, insufficient

vaccine take-up, because it is not consistent with relatively low infection rates among children, and because in heterogeneity analysis, we observe no protective effect of attending a middle school even with the highest vaccination rate.

Our study contributes to a literature investigating health externalities of vaccines, which has mostly focused on the flu vaccine. In an early experiment, Carmen et al. (2000) find that offering vaccines to long-term care hospital staff reduced elderly patients' mortality. Ward (2014) and White (2021) find influenza vaccinations among the non-elderly and among healthcare workers improve health outcomes among the elderly. Carpenter and Lawler (2019) find that adolescent TDap booster mandates reduce pertussis incidence among 0-4 year-olds.² We add to this literature in two ways. First, we provide quasi-experimental evidence on the direct effects of the COVID vaccine. Second, our finding that the COVID vaccines protect family members but not schoolmates provides the first evidence on the reach of the vaccine externality. This evidence is important because it suggests that strong direct effects and even household level indirect effects need not imply large indirect effects in contexts with moderate interaction, weakening the case for mandates in some settings.

We also contribute to research on pandemic mitigation. Previous work examines mobility and economic consequences of non-pharmaceutical interventions (Chetty et al., 2020; Kong and Prinz, 2020; Goolsbee and Syverson, 2021; Alexander and Karger, 2021; Gupta et al., 2021; Cronin and Evans, 2021). Our work is closest to research investigating health consequences of interventions such as masking (Abaluck et al., 2021; Ginther and Zambrana, 2021), college campus closure policy (Andersen et al., 2022) and shelter-in-place orders (Dave et al., 2021; Berry et al., 2021; Friedson et al., 2021). Especially relevant, Acton et al. (2022) show that college vaccine mandates reduce local COVID incidence and mortality, even among people too old to be college students. They show that college vaccine mandates had important positive health effects but do not disentangle direct and indirect effects. We complement their work by showing that vaccines' direct effects are large and indirect effects are concentrated among close contacts,

²We highlight papers that use experiments or quasi-experiments to estimate externalities. Other papers estimate effects of vaccines on secondary attack rates, but do not address selection-on-unobservables into vaccination. See Halloran et al. (2003) on pertussis and Lyngse et al. (2022) on COVID.

appearing in households but not schools.

2 Background and data

2.1 Vaccines can prevent illness without preventing infection

Clinical trials establish the in-sample (direct) effect of vaccines on illness, hospitalization, and death. But they do not test whether the vaccines reduce transmission, as this is not required by FDA. Preventing transmission requires “sterilizing immunity,” meaning that the vaccine prevents the virus from entering cells and replicating itself in the vaccinated individual. Most vaccines do not produce complete sterilizing immunity (Caddy, 2021). For example, the rotavirus vaccine and the Hepatitis B vaccine protect against illness *without conferring sterilizing immunity*. In contrast, the smallpox and measles vaccines do produce sterilizing immunity (Baker et al., 2019; Werner et al., 2013). Thus, even vaccines highly effective against disease need not prevent circulation. The reasons for this are complicated and contextual. For example, viruses may circulate by colonizing nasal passages without body-wide infection, and some vaccines may be less effective at preventing such local colonization (Bleier et al., 2021). Some partially vaccinated patients in the Moderna trial appear to have experienced such localized infections (Creech, 2022). In practice, the degree of sterilizing immunity may depend on multiple factors: recently vaccinated people are likely less transmissible, and vaccines are likely less effective in preventing transmission of variants than of the original virus for which they were designed.

2.2 Rolling out the COVID Vaccines

The COVID³ vaccines were approved in late Fall 2020. In the months after approval, demand exceeded supply, so policy makers used eligibility criteria to ration the vaccine. Most states established eligibility groups based on a combination of age, occupation, and risks. Eligibility was expanded one group at a time until everyone for whom the vaccine was approved was also eligible. Our focal state, Indiana, granted

³Throughout, we use the term “COVID” to refer to both SARS-CoV-2, the virus, and COVID-19, the disease.

eligibility to health care workers, first responders, and nursing home residents on 14 December 2020. In contrast to most states, however, subsequent eligibility groups in Indiana were determined only by age. Appendix Figure A.1 shows the date of eligibility, by birth cohort, in Indiana. Most age groups became eligible in short succession, with delays of 22 days or less between age groups, until 12-16 year-olds became eligible on May 12. Five-11 year-olds did not become eligible until November 3.⁴

Two aspects of the vaccine roll out are especially relevant for our empirical analysis. First, older people become eligible before younger people. Because age is a powerful risk factor, the age based rollout means that the riskiest people became eligible soonest and became vaccinated soonest. This makes it difficult to study how vaccination effects spill over to unvaccinated, high-risk family members. Second, most delays in eligibility between consecutive age groups are short, which makes it difficult to use eligibility-based variation to learn the effect of vaccination on COVID infection. The challenge is that take-up is neither immediate nor complete, and conditional on take-up, the immune response does not appear for 7-14 days. Thus there is likely a several week lag between eligibility and infection reduction.

For this reason, we focus on ages 11 and 12, where there was nearly a six month gap in eligibility. This relatively long gap provides the basis for several different identification strategies. We explain each strategy in detail below. At a high level, however, our approach is straightforward. We measure *direct effects* by comparing the outcomes of 12 and 11 year-olds themselves. We measure *indirect effects* at the household level by comparing adults living with 12 year-old vs 11 year-old family members, and at the school level by comparing people with 12 year-old vs 11-year old schoolmates.

2.3 Vaccines, mitigation, and behavioral responses

The roll out of the vaccine occurred in the context of ongoing pandemic mitigation strategies. Mitigation efforts include social distancing, working from home, and school policies such as remote instruction, mask mandates, and staggered class times to reduce mixing among students. It is possible that the

⁴For older ages, three vaccines were available (Moderna, Pfizer, and Johnson and Johnson), but for 16 year-olds and younger, the Pfizer vaccine was approved first. All our estimates of direct and indirect effects therefore pertain to the Pfizer vaccine.

widespread availability of the vaccine induced risk-altering behavioral changes – i.e. Peltzman (1975) effects. Both ongoing mitigation and behavioral responses influence vaccine effectiveness. Our estimates of direct and indirect effectiveness are inclusive of any behavioral responses and are estimated during a time when mitigation efforts remained in place. Because they are inclusive of possible behavioral changes, they are informative about the effectiveness of the COVID vaccine in slowing the pandemic. Policy decisions on vaccinations depend on how vaccines effect the spread of disease, given existing mitigation and subsequent behavioral responses to the vaccines—exactly the direct and indirect effects we estimate.

2.4 Electronic Medical Records and Registry Databases

We obtained data from two databases maintained by Regenstrief Data Services (Regenstrief Institute, 2022). The first is the Indiana Network for Patient Care (INPC) database, which consists of encounter and other medical records from health care providers throughout the state. The INPC was created to facilitate data sharing among Indianapolis hospitals, but membership has grown so that nearly all inpatient hospitals in the state participate. Several large health systems also contribute data, meaning much of the state’s healthcare providers are also included in the data. The second database is a registry of nearly all COVID lab (PCR) tests and COVID vaccinations conducted in the state of Indiana. This database grew out of a contract between the Regenstrief Institute and the state to develop a COVID dashboard. Analysts at Regenstrief Data Services generated patient identifiers to link individual records from these two databases.

The records provide basic demographic information along with data on health care encounters that can be used to construct measures of health care utilization, including hospital ER and inpatient admissions. INPC’s coverage is quite broad: the full data set contains 6.3 million unique person identifiers, while the state of Indiana has about 6.2 million residents.⁵

⁵There are more people in our sample than in the state of Indiana because the full data set stretches back to 2004, and so our data include out-migrants and decedents. It also contains records for out-of-state residents who obtain care in-state, although this appears rare in practice. Finally it is possible that the same person could have multiple IDs because of, for example, a name change.

2.5 Addresses, Household Assignment, and School Assignment

Two of our research designs require us to know what household a child lives in and what school they likely attend. We construct household and school assignments based on home location information derived from health care encounters. At each encounter, patients report an address, and this address may change over time.⁶ Our research using address information is governed by a data security plan, and was approved by the Indiana University Institutional Review Board. These address records form an “address log” of non-uniformly spaced, non-standardised text strings, containing the addresses from each encounter, along with the encounter date. In a pre-processing step on a non-networked computer (to maintain confidentiality), we work with the address log to impute some geographic information. We clean the text strings in the address log and then geocode each address. We successfully geocode 88 percent of addresses, meaning we match them to a latitude or longitude.⁶ This match rate is similar to the match rate in Sacarny et al. (2022) in the context of Medicaid enrollment addresses. We link each geocoded address to its schools⁷ (and in particular, sixth grade serving schools) using school catchment area maps provided by National Center for Education Statistics (2022b), derived from a survey of school districts ascertaining individual school catchment area boundaries. As not all school districts responded to the survey, 17 percent of addresses are not assigned to a sixth grade school. In these cases we assigned students to the nearest school, using school addresses from National Center for Education Statistics (2022a).⁸ After imputing school assignment, we drop the geocoded information from our analysis data set, but retain an encrypted address with linked school information.

In analyses requiring household or school assignment, we rely on a “primary address”, which we construct from the encrypted address information.⁹ We define the primary address as the address observed in the greatest number of years in 2017 and beyond. If there is no unique such address, we take the most

⁶We fail to geocode addresses either because they lack a location, as with a P.O. Box, or because we cannot map the location to an address.

⁷A given address can be assigned to multiple schools serving a given grade because some school districts allow school choice, so the school attendance boundaries overlap. We explain how we deal with this complication in Section 5.

⁸A remaining 3.7 percent of addresses are not assigned to any sixth grade school, because neither the closest elementary school nor the closest middle school serve sixth graders. We exclude these addresses from our school level analysis.

⁹We work with the first line of addresses only. However, 98.6 percent of addresses have a single line.

recent address among the qualifying addresses, and if that address is not unique, we break remaining ties at random. We focus on post-2017 addresses to reduce the chances of inconsistently assigning household members to different addresses, e.g. from before and after a move. As we describe in more detail in sections 4 and 5, we assign households and schools based on these geocoded primary addresses. For our household indirect effect analysis, we define a household as a collection of patients who share a primary address.¹⁰

2.6 Measures

We measure most outcomes at the person-month level. We define each person’s cumulative vaccination status using a dummy variable indicating whether the person has ever been vaccinated with both a first and second dose by the end of a given month or sooner. The measure is cumulative in the sense that a person stays vaccinated in every month after they first become vaccinated. Although our main analysis focuses on the effect of the second dose, in practice 94 percent of people with a first dose receive a full course (i.e. a second dose or one dose in a single dose regime).

We use the test registry data to construct a dummy variable indicating whether the person had a positive COVID PCR test during the month. We refer to these lab-confirmed COVID cases as COVID incidence. While our lab-confirmed case measure may miss COVID cases identified by rapid tests, such cases are likely rare during our period when rapid tests were often unavailable (Leonhardt, 2021). Appendix C shows that differential selection into testing is also not an important confounder, given the similarity of positivity rates we observe before and after the vaccine eligibility dates. Because we measure lab-confirmed COVID, we will miss asymptomatic cases and mild cases which do not prompt testing. Thus we interpret our outcome as measuring moderate and severe COVID cases rather than all COVID cases.

As a measure of severe COVID, we also form indicators for “hospitalized with COVID” and “emergency department visits with COVID.” We define these variables as a hospital admission or emergency

¹⁰Sacarny et al. (2022) face a similar challenge in the context of identifying household spillovers in Medicaid take-up, and they also infer households from addresses.

department visit within two weeks of a positive COVID test. In some analyses we also study *non*-COVID emergency department visits. We define non-COVID visits as those in which the patient did not have a PCR COVID test (positive or negative) in the 4 days before or 5 days after the visit (positive or negative).¹¹ Note that these definitions do not categorize all emergency department visits; ones with a negative COVID test are neither severe nor non-COVID. We define the variables this way because we want to use non-COVID visits as a coarse proxy for adverse vaccine events, and so they should be unrelated to COVID, including the suspicion of COVID.

In some analyses we use an individual-level risk score variable, meant to capture the risk that a COVID case becomes severe. To construct the risk score, we estimate a logistic regression for the likelihood of hospitalization-with-COVID, conditional on a COVID diagnosis in the prior two weeks. The predictors are age (in 5-year bins), demographics (race-ethnicity and sex) and comorbidities. The comorbidities are those used in Ellen et al. (2023) and align with the risk factors identified by the CDC that are straightforward to measure in claims data (Center for Disease Control and Prevention, 2023). In estimating the model, we use diagnoses measured in the year ending February 28, 2020 (i.e. the year prior to COVID onset), and we use hospitalizations during 2020 only, prior to vaccine availability. The risk score is the predicted probability of hospitalization with COVID, given diagnoses measured in 2020. Because we estimate the risk score using hospitalizations prior to vaccine availability, this measures the risk that a COVID infection becomes a severe case, in the absence of vaccination.

3 Direct Effects of the Vaccine

We begin by estimating the direct effects of the COVID vaccine, in particular the Pfizer vaccine because it was the first to be approved for people under age 17. Direct effects are driven by the causal effect of vaccination on COVID infection among people who are actually vaccinated. These direct effects are likely a necessary condition for indirect effects, but existing evidence on them is limited.

¹¹This procedure follows the approach used in Sacks et al. (2022), who show that the hazard of testing is especially elevated in this time window around hospitalizations.

3.1 Research Design - Direct Effects

To identify the direct effects of vaccines on the vaccinated, we construct a research design that exploits date-of-birth based eligibility criteria. The treatment group consists of people born in the year ending 12 May 2009, all of whom became eligible for the vaccine on 12 May 2021. The control group consists of people born in the year *beginning* 3 November 2009, roughly six months after the last treatment group birthday. To build our analysis sample, we extract the $N = 133,013$ patients in the Regenstrief databases born in these ranges and alive as of January 1, 2020, then measure monthly outcomes using the vaccine and testing registries, forming a balanced panel.¹² The key difference between the treatment and control group is their date of vaccine eligibility: the control group did not become eligible for the vaccine until 3 November, six months after the treatment group. However, the summary statistics in Table 1 show that the two groups are otherwise similar in their demographics, although pre-period COVID incidence is slightly higher in treatment than control.

We estimate the effect of earlier vaccine eligibility with the canonical regression implementation of the difference-in-difference estimator:

$$y_{it} = \beta_0 + \beta_1 \text{EarlyElig}_i + \beta_4 \text{post}_t + \beta_3 \text{EarlyElig}_i \cdot \text{post}_t + \epsilon_{it}. \quad (1)$$

The outcome y_{it} is an indicator for vaccination status, positive COVID status, or non-COVID emergency department visit (a measure of adverse events) for person i in month t . $\text{EarlyElig}_i = 1$ for students that belong to the early eligibility birth cohort born between 13 May 2008 and 12 May 2009, i.e. the treatment group. post_t indicates the post-June-2021 months. We begin the sample in March, 2020, as there is essentially no COVID incidence before then, and end in December, 2021. We include November-December of 2021 even though the control group became eligible then because, as we show below, take-up is low in the control group at this time, so these additional months improve power. We estimate standard errors using a cluster robust variance matrix that allows for dependence from repeated observations on the same individuals.

¹²Appendix Table A.1 shows how our sample size changes as we impose our inclusion criteria, for this and all designs.

The difference-in-difference design is robust to two important sources of possible confounding. The coefficient on $EarlyElig_i$ should capture any time invariant differences—including birth year and the pre-period difference in COVID-incidence—between the treated and control birth cohorts that may affect COVID related health outcomes. Similarly, the coefficient on $post_t$ should capture any cohort invariant differences in epidemiological conditions between 2020 and 2021. The coefficient on the $EarlyElig_i \cdot post_t$ interaction term represents the average effect of vaccine eligibility on the treated cohort, under the usual parallel trends and no-anticipation assumptions. One potential threat to the validity of this design is indirect effects. Treated children may have more vaccinated peers—friends and schoolmates, for example—than do untreated children, because of age-based vaccine eligibility. Our estimates here potentially reflect a mixture of direct and indirect effects, and hence potentially overstate the direct effects, though they are likely valid for testing the null hypothesis of no direct effect. In subsequent sections, however, we estimate small and insignificant school- and grade-level indirect effects, suggesting little bias in the direct effect estimates from indirect effects from peers.

Although vaccine eligibility in 2021 differs between our treatment and control group, many eligible children do not end up vaccinated. As a result, the difference in difference specification provides reduced form—intent to treat—estimates of the effects of vaccination. To isolate the effects of actually being vaccinated, we also estimate instrumental variables models in which we instrument for vaccination take-up using the interaction $EarlyElig_i \cdot post_t$. The second stage equation is

$$y_{it} = \gamma_0 + \gamma_1 EarlyElig_i + \gamma_2 post_t + \gamma_3 \widehat{vaccinated}_{it} + \nu_{it}. \quad (2)$$

In the model, γ_3 gives the percentage point effect of vaccination on outcome y_{it} for compliers. In Appendix B we show how to translate this to an estimate of the more familiar vaccine effectiveness metric.

3.2 Results - direct effects

Our key results on the direct effects are evident in the raw time series of vaccination rates, COVID incidence, and non-COVID emergency room visits for the early eligible treatment and late eligible con-

control groups, plotted in Figure 1. The figure shows that vaccine eligibility increases vaccinations, reduces COVID incidence, and has no discernible effect on non-COVID emergency room visits. Starting from the top panel we see that vaccine take-up grows steadily for the treatment group when they become eligible, with no vaccination in the control group until their eligibility date six months later. The middle panel shows that, in the pre-period, the treatment and control groups had essentially equal COVID incidence, suggesting little or no confounding. The treatment group experiences a slightly larger increase in incidence during the Winter 2020 peak than does the control group; this differential seasonality would if anything attenuate our estimates of protection. After becoming vaccine eligible, the treatment group diverges from the control group; the vaccine-eligible students in the treatment group experienced lower COVID infection rates in each of the last four months of 2021.¹³ The final panel shows that vaccine eligibility has no apparent effect on adverse events, measured here as non-COVID emergency room visits. Treatment and control show nearly identical levels and trends throughout the sample period, with no divergence after vaccine eligibility. For effects on all-cause and COVID-related visits, see Appendix Table A.2.¹⁴

We report DID and IV estimates in Table 2. The first two columns show that vaccine eligibility increases the vaccination rate by 23 percentage points, with nearly identical impacts on first and second doses. The next column shows that early eligibility reduces COVID incidence by 0.3 percentage points. Our instrumental variables estimates, which adjust the reduced form results to account for vaccinations induced by eligibility, indicate that vaccination itself reduces COVID incidence by about 1.3 percentage points, for a complier vaccine effectiveness of about 80 percent. (See Appendix B for details on vaccine effectiveness.) This estimate is roughly comparable to, but somewhat smaller, than the 95% effectiveness reported in the clinical trials for the mRNA vaccines (Baden et al., 2020; Polack et al., 2020). The results in column 4 imply that earlier eligibility has no effect on non-COVID emergency room visits (confidence

¹³This divergence does not occur until September 2021, four months after initial vaccine eligibility. This delay is unsurprising: vaccine take-up grew over time, vaccines take time to generate an immune response, and COVID prevalence was fairly low in May-July of 2021, but grew dramatically in August and September as the Delta variant circulated and school resumed.

¹⁴A particularly critical adverse event is “multi-system inflammatory condition.” We observe zero cases of this in either treatment or control.

intervals rule out effects larger than about +0.5 percentage points), suggesting that the vaccine did not lead to widespread, severe adverse side effects. These results are robust to alternative sample inclusion criteria such as a fuller set of months, or looking at a wider range of birthdays (Appendix Table A.3).

One potential concern with this estimate is the higher treatment group COVID rate in the pre-period. This could cause violations of the parallel trends assumption if the higher COVID rate generated long-lasting natural immunity, reducing treatment group COVID incidence in the post period, relative to the control group. In practice this type of confounding is unlikely, for two reasons. First, the difference in pre-period COVID rates is only 0.08 percentage points, small relative to the 0.3 percentage point differential decline in COVID incidence. Second, in the fall of 2021, the primary variant was the Delta version, against which ancestral variant antibodies are less effective (Mlcochova et al., 2021).

4 Household-wide Indirect Effects

The evidence so far suggests that Indiana's age-based vaccine roll out generated exogenous variation in vaccine take up among 12 year olds relative to 11 year olds. The eligibility induced vaccinations appear to have reduced COVID infection rates by about 1.3 percentage points among the vaccinated 12 year olds, which implies that the vaccine is about 80 percent effective. These effects are very likely driven by the direct effect of the vaccine. But the vaccines could have important indirect effects if vaccination reduces COVID transmission rates from vaccinated to unvaccinated people. To the extent that indirect effects exist, they may matter the most in small peer groups that share physical spaces in ways that is easier for the virus to spread from one person to another. In this section of the paper, we study the indirect effects of the vaccine within households by exploiting the Indiana vaccine roll out and the age composition of individual households.

4.1 Research Design - Household indirect effects

Our goal is to identify household level indirect effects of vaccination. An indirect effect is a spillover or peer effect in which vaccinating one person confers health benefits on a different person. For example, if

persons i and j live in the same household, an indirect would imply that vaccinating person i would reduce the probability that person j will contract COVID. An ideal experiment for estimating the indirect effect of the vaccine might randomly assign people to a treatment group and a control group. In the treatment group, we would vaccinate one of the other people in each treatment group member's household. Comparing downstream COVID infection rates among the focal people in the treatment and control groups would identify the causal effect of the household vaccination rate on a person's own COVID infection rate. In other words, identifying household-level indirect effects requires exogenous variation in one household member's vaccination, independent of vaccination and COVID risk of other members.

We approximate the ideal experiment by comparing households that are differentially affected by Indiana's age-based vaccine roll out because of the age composition of the children in the household. The basic idea is to compare people living in households with a child born between May 2008 and May 2009 to households with a child born between November 2009 and November 2010. Children born in the May 2008-May 2009 cohort were eligible for the vaccine in May 2021, six months sooner than children in the November 2009-November 2010 cohort. As a result, households with the May 2008-May 2009 children experienced higher household vaccine eligibility and higher household vaccination rates in the summer and fall of 2021 than the households with children born a year later, despite similar vaccination rates and COVID rates among other household members.

We operationalize this idea using a difference-in-difference design. To build the analytic sample, we start by defining children born between May 2008 and May 2009 as "early eligibility children". Children born between November 2009 and November 2010 are our "control children", as in the direct effects design. But the treated and control children are not the subjects of our study of indirect effects. Instead, the *treatment group* consists of people age 30 and older who live with at least one early eligible treated child, no control child, and no "interim" children born between June and October 2009. The *control group* consists of people age 30 and older who live with at least one late eligible control child, no treated child, and no interim children. These inclusion criteria help ensure the key difference between the treatment group and control group is their exposure to a child with earlier vs later vaccine eligibility. We further

restrict the sample to people alive as of January 1, 2020, with 8 or fewer household members to help with measurement error in household assignment, as explained below. Appendix Table A.1, panel B, shows how our sample size changes as we impose the inclusion criteria.

We build a balanced panel of the treatment group and control group members that runs from March 2020 to December 2021. The post period is June 2021 through December 2021, and the pre-period is March 2020 through May 2021. With these alterations in sample inclusion criteria and date ranges, we again estimate difference-in-difference regressions. The coefficient here is an estimate of the effect of increasing the number of household members that are eligible for the vaccine. We view this as a reduced form or “intent to treat” estimate of the effects of an exogenous increase in the vaccination rate among other household members on the health of the adults.

The difference-in-difference design relies on the parallel trend assumption that—in the absence of differential household vaccine eligibility—the COVID infection rates in the treatment and control groups would move in parallel. An important way this condition might fail is that treatment households may consist of older adults relative to those in control households because treatment households have older children. Older adults are vaccinated sooner, independently affecting COVID trends. We address this challenge by re-weighting the control group to match the age distribution of the treatment group.¹⁵ We show below that, after re-weighting, pre-period trends are parallel and pre-period outcomes are quite close on average.

Given incomplete take-up, the intent-to-treat estimates from the DID models can be difficult to interpret or compare to other estimates, such as the effect of own vaccination. To quantify the spillover more precisely, we also estimate two stage least squares regressions for the effect of household vaccination:

$$Y_{it} = \beta_0 + \beta_1 \widehat{Vacc}_{-it} + \beta_2 treat_i + \beta_3 post_t + \epsilon_{it} \quad (3)$$

$$Vacc_{-it} = \alpha_0 + \alpha_1 EarlyEligChildPresent_i \cdot post_t + \alpha_2 EarlyEligChildPresent_i + \alpha_3 post_t + \epsilon_{it} \quad (4)$$

¹⁵We implement the re-weighting by estimating a propensity score for treatment, modelling the propensity score as a function of fixed effects for each age. Thus our re-weighting is equivalent to (many-to-one) propensity score matching.

In the model, $Vacc_{-it}$ is the vaccination rate of all members in i 's household except i herself, and the excluded instrument is the interaction $EarlyEligChild_i \cdot post_t$. Our primary outcome— Y_{it} —is a dummy variable indicating that the person had a lab-confirmed COVID infection during month t . We also consider emergency department and hospitalizations with COVID. The first stage difference-in-difference regression measures the effect of the early eligible child on the vaccination rate in the household where the treated adult is living. The $EarlyEligChildPresent_i \cdot post_t$ interaction term serves as the instrumental variable in the second stage equation. The instrumental variable estimator isolates the average causal response to the increased vaccination rate among the complier households. We describe characteristics of complier households below.

While Equation 3 appears to impose unrealistic linearity assumptions, in fact β_1 can be interpreted through the lens of the average causal response theorem developed in Angrist and Imbens (1995). In particular, in a model with unrestricted causal effects of marginal increases in vaccination rates, β_1 is (100 times) the weighted average marginal effect of a 1 percentage point increase in the household vaccination rate, averaging over all possible rates (0, 0.01, etc.), with weights proportional to the share of compliers whom the instrument brings from below that rate to above it. (See Appendix D.) Below we calculate these weights to see which levels of vaccination are reflected in our estimates.

4.2 Measurement Error

Because we do not perfectly observe household identifiers, both the endogenous regressor ($vacc_{-it}$) and the instrument ($treat_i \cdot post_t$) may be measured with (potentially non-classical) error. To understand the issue, imagine that household 1 consists of a 12 year-old treated child and two adults, and household 2 consists of a 9 year-old child and two adults. Household 1 lives at an address until 2020, when household 2 moves in. Our procedure may incorrectly lump household 1 and household 2 together, misclassifying the treatment status of household 2. Misclassifying the household membership may also lead to mis-measured household vaccination rates, because the household 1 child is potentially vaccinated since she is eligible. More generally, when we incorrectly classify a household as treated, we are likely to overstate

its vaccination rate, and when we incorrectly classify it as control, we are likely to understate its vaccinate rate.

The consequence of this measurement error is that the structural parameter β_1 , the effect of household vaccination on own COVID, is likely *understated*. To see why, note that the measurement error in our endogenous regressor is positively correlated with the misclassification in treatment status, inducing a positive bias in the first stage. The measurement error leads to a downward attenuation bias in our reduced form since the measurement error in treatment status is likely independent of COVID status. Thus the IV estimate is too small. This is important because our basic approach, which does not fully correct for measurement error, ends up finding large and significant spillovers.

Beyond noting that the bias here works against our ability to detect a household spillover, we address the measurement error in three ways. First, in our primary analyses we limit the sample to households with 8 or fewer members, as in Sacarny et al. (2022). Doing so helps exclude addresses that are erroneously assigned to multiple families. Second, in a robustness check, we take advantage of multiple years of address data—and hence multiple measures—to estimate obviously related instrumental variables models (ORIV, Gillen et al. (2019), described in detail below below). These provide noisier but unbiased estimates of household spillovers. The estimates are quite similar to our baseline estimates, again implying low measurement errors. Third, although we cannot directly quantify the measurement error, below we report *adult* vaccination rate among “never taker”, “complier”, and “always taker” households, i.e. households where children are never vaccinated, vaccinated if eligible, and vaccinated even when ineligible. We find much higher adult vaccination rates in complier and always taker households than in never taker households. Importantly, we use no information on adult vaccination rates to determine these categories, so this shows a strong correlation within household between child and adult vaccination rates—as we would expect if there was little measurement error in household assignment.

4.3 Summary Statistics

We report summary statistics for the household analysis sample in Table 3. The first column reports means for the treatment group, the second column for the control group (after re-weighting), and the third column the difference in means. There are some significant differences between treatment and control but overall they appear well balanced. In particular in the treatment group the proportions female and Hispanic are significantly higher, but the magnitude is small. Importantly, the pre-period difference in COVID rates is small and insignificant, indicating little confounding. Treatment households also have slightly fewer members, but slightly more members aged 14-18 (who also become vaccine eligible in late spring 2021). In principle this could mean that any difference between treatment and control in the post period is due not just to the treated child's vaccination but also to their older siblings'. Such additional vaccination is not a problem for our exclusion restriction, which requires that the only reason that a vaccine-eligible child affects parental COVID is by changing the household vaccination rate. It is also not likely to be important in practice, given the quantitatively small difference (i.e 0.013 people in the older age range, vs. an additional 1.00 person aged 12-13).

4.4 Results

We begin by showing event study plots in Figure 2. The top left panel shows the household vaccination rate in each month in the treatment group and in the age-weighted control group, along with a 95 percent confidence interval based on standard errors clustered on household. For both treatment and control adults, household vaccination rates increase steadily and equally from 0 to 15 percent throughout spring 2021 as more adults became eligible for the vaccine. But the two groups diverge in June 2021, when treatment group children become eligible for the vaccine. As the event study in the top right panel shows, the treatment-control difference grows by about 6 percentage points in June and continues growing to about 11 percentage points by September, where it levels off.

Importantly, the differential increase in household vaccination rate is not accompanied by an increase in own vaccination rates, as the graphs in the middle row of Figure 2 show. Own vaccination rates move

perfectly in parallel both before and after the adolescents in the treated household become eligible for the vaccine. We do however find spillovers to other children in the household, consistent with Humlum et al. (2022); Carpenter and Lawler (2019)). In Appendix Table A.4, we decompose the impact of a vaccine-eligible child on the household vaccination rate into vaccine take-up among focal children (i.e. treated or control children, aged 11-13), older children, younger children, and other adults. We do so by creating age-specific versions of the household vaccination rate, for example looking at vaccinations among 11-13 year-old children (as a share of the household size). About 90 percent of the overall effect comes from increased take-up among treated children, with small spillovers to older children. These cross-sibling spillovers are qualitatively consistent with

The bottom row of graphs in Figure 2 provides evidence on these indirect effects. The left panel shows that COVID infection rates in the treatment and control group move in parallel in 2020, and indeed have nearly identical levels of infection as well. Both groups experience a substantial increase in infection rates during the October-December 2020 seasonal peak. In the fall of 2021, when both COVID prevalence and adolescent vaccination rates are high, the COVID infection rates in the two groups diverge.

We quantify the reduced form impact of adolescent eligibility on adult outcomes in the first row of Table 4. Each column shows the DID estimate for a different outcome. Adolescent eligibility increases household vaccination rates by 11 percentage points, a highly significant estimate. It has no effect on own vaccination among adults. Our key result in column (3) is a statistically significant reduction in adult COVID incidence of about 0.1 percentage points per month, which is between 5 and 10 percent of the control group incidence during this time period. The difference-in-difference estimates mirror the graphical evidence in Figure 2, showing that COVID incidence fell in treated households relative to control households following adolescent vaccine eligibility. The results provide evidence against the null hypothesis of no household-level spillovers. However, when we look in column (4) at “severe COVID”, defined as an ER or inpatient admission with a positive COVID test in the prior 14 days, we find small and insignificant effects. Given the modest first stage, it is difficult to directly interpret these reduced form estimates quantitatively. We therefore turn to the two-stage least squares estimates.

The second row of Table 4 reports our estimate of β_1 , which implies that increasing a person’s (other) household vaccination rate from 0% vaccinated to 100% vaccinated reduces the person’s COVID infection rate by 0.7 percentage points. Under homogenous effects, this is the effect of fully vaccinating the rest of the household on an adult’s COVID incidence. Although unrealistic, the homogenous effects assumption makes it straightforward to compare the household-level spillover to the direct effect. We find a spillover among adults of about -0.7 percentage points for fully vaccinating the household, and a direct effect of about -1.3 percentage points among adolescents. Thus the indirect effects are about half as large as the direct effects. Put differently, the model suggests that for each COVID case prevented by the direct effects of vaccinating a 12 year-olds, half a case is prevented among household members.¹⁶

The homogeneity assumption is strong, however. Recall that our estimate of -0.7 can be interpreted as (100 times) the weighted average marginal effect of a 1 percentage point increase in household vaccination rate, with weights on each vaccination rate proportional to the share of compliers whom the instrument brings from below that rate to above it. (Appendix D.) We plot the distribution of these weights in Appendix Figure A.2. There is positive weight on all values of household vaccination, but the mass of weight is concentrated between about 33 percent and 75 percent, with especially little weight below about 25 percent. We therefore learn about the effect of increasing household vaccination at intermediate levels, and learn less about impacts at low household vaccination rates.

While the instrument therefore mainly reflects marginal household vaccination rates effects at intermediate rates, the adults in the complier households are themselves quite highly vaccinated. They live with other unvaccinated children. To see this, define treatment status by the December, 2021 vaccination status of the focal child, a binary treatment. That is, “always takers” are households that vaccinate their child by December 2021 regardless of early eligibility, never takers do not vaccinate, and compliers vaccinate their child only with early eligibility. Using standard tools for calculating complier characteristics, we find that the *adult* vaccination rate among complier households is 80 percent in December 2021, versus

¹⁶To see this more explicitly, note that vaccinating a 12 year-old raises the household vaccination rate by $1/(N - 1)$, where N is the number of people in the household, reducing COVID incidence among each of the $(N - 1)$ household members by $-0.007/(N_1)$ in expectation, producing a cumulative decline of -0.007 . As the vaccinated 12 year-old’s incidence declines by .013 in expectation, we have approximately an additional 0.5 cases prevented.

a never taker vaccination rate of 36 percent.¹⁷ Thus our household-spillover estimates reflect the effect of household vaccination among adults who are mostly vaccinated themselves and who live in households with high vaccination rates.

The high complier vaccination rate helps explain why we find no protective effects on severe COVID, defined as inpatient or ED visits with COVID, as we show in column (4). The high household vaccination rate among compliers implies that adults in complier households are quite likely to be vaccinated themselves. As vaccines are particularly effective at preventing severe cases of COVID, there is relatively little scope for transmission from children to adults to lead to severe COVID among our compliers.

In Table 5, we investigate heterogeneous effects along two natural dimensions: COVID risk and vaccine take-up. We first split the sample by risk score, which is a measure of a person's likelihood of hospitalization given COVID infection, estimating separate models for adults with above and below median risk scores. While the first stages are similar, we find larger spillovers for low risk score adults than for higher risk score adults.¹⁸ A possible explanation for this difference is that the highest risk group is more likely to be vaccinated themselves (55 percent own vaccination, vs. 46 percent among the lower risk group), and may be more cautious in other respect as well. We see a similar pattern when we split by age: for older adults, the protective effects of household members' vaccination is smaller, but own vaccination rates are higher. In no case do we see statistically significant impacts on severe COVID. In Appendix Table A.5, Panel A, we examine indirect effects in a separate sample: children aged 10 and younger, a population with a very low vaccination rate. Here we find a slightly smaller first stage and larger reduced form, which together imply a statistically significant indirect effect of -0.018.

In the final panel we split the sample on baseline vaccination status, defined as the vaccination status of the index adult as of May 2021. The first stage is much stronger among adults who are themselves vaccinated, and we find a significant effect of household vaccination for this group. For the unvaccinated

¹⁷Adult always-takers, who live in households where children would be vaccinated by December, 2021 even when the child is eligible late, constitute 2.6 percent of the sample, and are themselves likely to be vaccinated early. By December 2021, 85 percent of always takers are vaccinated.

¹⁸There are more adults with median-and-below risk scores than with above median risk scores because the risk score is fairly coarse.

sample, the first stage is relatively weak and the IV estimate is sufficiently imprecise that we cannot rule out either a null effect or an effect equal to the spillover for vaccinated adults.

We emphasize that, while we do not find that vaccinating one additional household member prevents severe COVID among others, our results show that vaccines reduce within-household transmission. In households with low vaccination rates, this reduced transmission could have important health spillovers, consistent with the relatively large effect we see among young children. In other contexts, such as health care workers interacting with patients, indirect effects might be large (e.g. White (2021)).

Comparison to prior findings Our findings of substantial household-level indirect effects of the COVID-19 vaccine are consistent with other work finding infection-reducing externalities from vaccinations. Ward (2014) shows that a universal flu vaccination campaign in Ontario increased vaccine take-up of the non-elderly by about 11 percentage points and reduced flu hospitalizations among the elderly, with no effect on non-elderly hospitalizations; similarly, White (2021) finds that increased non-elderly flu vaccinations in the US reduce elderly mortality. Carpenter and Lawler (2019) find that TDap mandate for middle schoolers increased their take-up by 13.5 percentage points, and reduced pertussis incidence by similar amounts among 0-4 year-olds as among middle-school aged children. These papers, using aggregate data, are unable to analyze the site of reduced transmission. Our results suggest that household-level transmission is an important component. White (2021) also shows that vaccination of health care workers prevents community spread, showing that not all prevented transmission is at the household level. Prior work has also found more substantial indirect health impacts, in the form of reduced hospitalizations for influenza. We do not find such impacts, likely because our sample exposed to indirect effects is relatively highly vaccinated and not at the greatest risk from severe infection, as it mainly includes parents rather than grandparents.

Sensitivity analysis We explore the sensitivity of our results to three key design choices: the focus on adults, the time period used to define treatment and control children, and the handling of measurement error.

In our main analysis we study indirect effects among adults rather than younger children in the house-

hold, for whom COVID is relatively less threatening, but vaccination rates are lower. In Appendix Table A.5, Panel A, we reproduce our main estimates but changing the sample so that it consists of children aged 0-10 (alive on January 1, 2020) and otherwise meeting our inclusion criteria. We find a statistically significant indirect effect, more than twice as large as our estimated indirect effect for adults. Thus the indirect effect we document for adults is also evident for children.

Our results are also not overly sensitive to the time period used to define treatment and control children. In our baseline analysis we use a 365 day window to define treatment and control children. In Panels B and C of Appendix Table A.5, we show results for two alternative approaches, using 180 or 540 day windows. The shorter cutoff results in a larger point estimate for the indirect effect, but also a larger standard error; the estimates are statistically indistinguishable.

The final panel shows that our results are also robust to a procedure which better handles measurement error: obviously related instrumental variables (Gillen et al., 2019). This procedure differs from our baseline in that we use addresses from one time period to measure treatment status and form the sample, and addresses from a separate time period to measure household vaccination rates. Using non-overlapping time periods makes it more likely that our measurement error is independent, and hence solved by instrumental variable methods. To implement ORIV, we start by imputing address characteristics based on the full set of residents ever recorded at each address. But we assign these patients to two primary addresses, separately using the address log from 2016-2020 and from 2021-2022. (We choose these ranges because there are roughly equal numbers of entries in each time period.) This gives us two measures, one from each time period, of instrument status, sample eligibility (i.e. household size and presence of treatment, control, and interim children), and household vaccination rate. In the ORIV approach, we stack the two data sets—one based on the first address and the other on the second address—and then use each time period once for the endogenous regressor and once for the instrument. We continue to cluster standard errors at the address level (as measured in the instrument), accounting for the dependence this stacking generates.

The ORIV estimates end up fairly similar to our baseline estimates. The first stage and reduced form

estimates are quite similar to our baseline estimates, and so is the IV estimate.¹⁹ This sensitivity analysis shows that our main results—that the vaccines do produce indirect effects on COVID infection within households—are robust to measurement error, and indeed the similarity of the point estimates suggest that there is little error in our main measures.

5 School-wide Indirect Effects

Our results to this point imply the vaccines do seem to have non-trivial spillover effects, at least inside households where people share important living spaces and are in close contact. However, households are not the only environment where indirect effects of vaccination may be epidemiologically and economically important. People also interact in environments such as work sites and schools. These locations are less intimate and congested than households, but they may also be places where vaccine spillovers could be important. In this section of the paper, we estimate indirect effects of the vaccine in schools.

5.1 Research Design – School indirect effects

Our empirical strategy to identify the effects of school-wide vaccination rates on individual COVID risk is based on a comparison of sixth graders in middle and elementary schools. The idea behind this comparison is that middle school students in the fall of 2021 are mostly eligible for the COVID vaccine, as all seventh grade and older students are eligible. In contrast, elementary school students are almost all ineligible. Because some sixth graders go to middle schools and others go to elementary schools, this difference in eligibility leads to large differences in school-wide vaccination rates. Our comparison of middle school and elementary school sixth graders therefore identifies the effect of having more vaccinated schoolmates, holding fixed both own vaccination rate and grade-wide vaccination rates. It provides evidence on arm’s length indirect effects.

We estimate the reduced form effect of exposure to vaccine-eligible peers using a difference-in-differences

¹⁹The point estimate is actually smaller, despite our argument of attenuation bias from measurement error. This difference is potentially explained by the different sample; we limit to adults with two household measures in the ORIV analysis.

regression, and student-month level data. Here treatment is measured by $middle_i$, an indicator for assigned to middle school, and our control group consists of sixth graders assigned to elementary schools. The post period is August-December 2021, and the pre-period is August-December 2020. We cluster standard errors at the school district level. The difference-in-difference model adjusts for time-invariant differences between middle and elementary schools, as well as for common trends in incidence and vaccination that could arise due to pandemic waves and time-varying eligibility.

To put the difference-in-difference-estimates on a scale comparable with the direct effect estimates and the household-level indirect estimates, we also estimate two-stage least squares models for the effect of school wide vaccination rates on COVID incidence:

$$Y_{it} = \beta_0 + \beta_1 \widehat{VaccSchool}_{it} + \beta_2 Middle_i + \beta_3 post_t + \epsilon_{it} \quad (5)$$

$$\widehat{VaccSchool}_{it} = \alpha_0 + \alpha_1 Middle_i \cdot post_t + \alpha_2 Middle_i + \alpha_3 post_t + \epsilon_{it} \quad (6)$$

Here the outcome is COVID incidence and our interest is in β_1 , the effect of a 100 percent increase in the school wide vaccination rate on student i , a measure of indirect effects (assuming, as we verify below, that middle school attendance influences COVID incidence among sixth graders by changing their school-wide vaccination rate but not their own vaccination rate). This model appears to assume, unrealistically, homogeneous effects of peer vaccination. However, the main goal of the 2SLS model is just to rescale the DID estimate so that it is comparable to our other estimates. As we argue in the case of household indirect effects above, results from Angrist and Imbens (1995) imply that β_1 can also be interpreted as the average effect (per unit) of increasing school vaccination rates from elementary school levels to middle school levels, roughly going from 5 percent to 25 percent.

5.2 Data: School Characteristics and Analysis Sample

Implementing our empirical strategy requires assigning sixth grader students to a school type: middle school or elementary school. To do so, we start by assigning addresses to schools using school catchment

area maps and, for addresses not covered by a catchment area, the closest school (as described in Section 2.3). Using school information from National Center for Education Statistics (2022a), we define each school that enrolls sixth graders as either “six and up” (e.g. grades 6-8), “six and down” (e.g. K-6), or “six between” (e.g. K-12). We define an address as a “middle school” address if *all* the sixth grade serving schools it maps to are “six and up”, and “elementary school” if all the sixth grade serving schools it maps to are “six and down.” If an address is assigned to both “six and up” and “six and down” schools (because of a school choice zone), then we exclude it from our main analysis sample. Thus we end up assigning students to school *types* rather than exact schools. This type assignment is all that is required for our empirical strategy, and is likely subject to less error than is assigning students to actual schools. This is because most school districts use a consistent middle school grade range, so as long as a student is in the right district, we have the right school type. The main source of measurement error is private school attendance. As about 8 percent of Indiana students attend private schools,²⁰ in the worst case this attenuates our estimates by about 16 percent (i.e. if we misclassified school type for all private school students).

We additionally require data on school-wide vaccination rates, and in robustness tests, the presence of mask mandates and in-person instruction, which might impact indirect effects. We define the monthly school-level vaccination rate as the monthly vaccination rate among all students whose address maps to that school. We measure mask mandates using data from Waldron (2022b), who collected data on school mask policies for most Indiana school districts through September, 2021, by visiting district websites and monitoring Google alerts (Waldron, 2022a). We extend the data through summer, 2022, by visiting each school districts’ website and checking for news releases about COVID mitigation policies. We use `archive.org` to fill in missing gaps where possible. We focus on mask mandates for students (as opposed to recommendations, or staff mandates). Our school instruction modality data come from COVID-19 School Data Hub (2022), who report instruction modality by school and week for the 2020-21 school year. We focus on in-person instruction as a measure of modality, and we assume all schools are in person for the 2021-22 year. These variables are measured at the school level, but some students are

²⁰See <https://www.in.gov/doe/about/news/indiana-k-12-school-enrollment-grows-for-2021-2022-school-year/>

assigned to multiple schools (if their address falls in multiple catchment areas). For such students we use the average vaccination rate, masking policy, or modality, averaging over the schools she could attend.

To construct our analysis sample, we begin by selecting records of sixth graders, with grades imputed assuming all students begin school in the year in which they are aged 5 years old on August 1, and progress one grade per year. We further limit the sample to students alive on January 1, 2020, with an address, assigned to a sixth grade school type, unambiguously middle or elementary school. In our primary sample we exclude students with imputed school assignment (because their district did not contribute data to National Center for Education Statistics (2022b)), though we show our results are not sensitive to this final restriction. Appendix Table A.1 shows how the sample size changes as we impose these restrictions.

We report summary statistics for the treatment and control group in Table 6. The treatment and control group appears well balanced on demographics, risk scores, and pre-period COVID rates.

5.3 Results – School indirect effects

In contrast to the protective direct effect and indirect effect at the family level, we find no protective indirect effect at the school level. Our key results are again evident in the simple trends, which we plot in Figure 3, for sixth graders in middle and elementary schools. Treatment sixth graders experience a large increase in the vaccination rate of their schoolmates, relative to control sixth graders, but no differential decrease in COVID-19 incidence. The top panel shows that vaccination rates are zero for both groups until May, 2021, when they diverge sharply. By fall 2021, treatment group sixth graders go to schools in which about one in five students are vaccinated, while for the control group the number is closer to one in 100.²¹ We further show in Appendix Figure A.3 that COVID rates among 7th graders fall, relative to 5th graders, after vaccine availability. Turning to COVID-19 incidence, in the middle panel, we see near identical levels and trends in the pre-period for treatment and control, suggesting little if any confounding. Incidence increases in fall 2021 for both groups. However, despite the large relative increase in schoolmate vaccination for the treatment group, and lower schoolmate COVID incidence, we see no

²¹Peer vaccination rates in the control group increase and then decrease steeply in May-July 2021, because sixth graders were vaccine eligible in the spring of 2021. This pattern does not influence our estimates because we report fall-on-fall differences.

relative decrease in COVID-19 incidence during this period. By the late fall, sixth graders in treatment and control alike become vaccine eligible themselves, and in the bottom panel we see some evidence of higher vaccine take-up among the treated sixth graders.

We report difference-in-difference estimates and standard errors, by time period, in Table 7. Going to a middle school induces a 23 percentage point increase in the school-wide vaccination rate; this difference persists into the late fall, when sixth graders become vaccine eligible, because take-up is fairly low but continues to grow among the older students. This school-wide vaccination increase is entirely driven by other grades: the impact on own vaccination (and hence sixth grade vaccination rate) is an insignificant 0.3 percentage points. We find no protective effect of the increase in schoolmate vaccination. The reduced form effect of attending a middle school when middle schoolers are vaccine-eligible is a 0.1 percentage point *increase* in COVID incidence. Scaling by the impact on school vaccination, our IV estimate implies that attending a fully vaccinated school would increase own COVID incidence by about half a percentage point. This estimate is significantly different from -0.0074 ($p=0.03$), the estimated family level indirect effect. (However, accounting for uncertainty, the point estimates are not significantly different from each other.)

Sensitivity analysis Our finding of no spillovers from vaccinated schoolmates is robust to alternative samples and specifications, as we show in Appendix Table A.6. In particular, we expand the sample to include students with imputed school assignment, limit the sample to include students with unique schools (rather than just unique school type), and control for instructional modality and mask mandate presence. Across these specifications the reduced form and instrumental variables estimates are positive and insignificant, and we continue to reject the null hypothesis that the estimated school-wide indirect effect equals -0.0081, the point estimate for the household indirect effect.

Comparison to prior findings As discussed in Section 4, prior work has found important indirect effects of vaccines using aggregate data (Ward, 2014; Carpenter and Lawler, 2019; White, 2021). Our null finding here is not inconsistent with this prior literature. Rather, combined with our earlier results, they suggest that indirect effects are likely to be context specific, differing at least between schools and house-

holds. There are at least two context-specific factors that influence the magnitude of indirect effects of a marginal vaccination: the level of vaccination (see, e.g., Goodkin-Gold et al. (2020)) and the extent of mixing between the marginally vaccinated person and the rest of the community.

The level of vaccination is a potentially important difference between the household-level analysis and the school-level analysis, as our natural experiment increased the vaccination rate to over 50 percent in the household analysis, but only to about 25 percent in the school level analysis. We investigate the importance of the vaccination rate in heterogeneity analysis. Specifically, we divide our treatment and control groups into quartiles of their middle school's school-wide vaccination rate (as of October, 2021), and re-estimate our DID and IV models within each quartile.²²

We show the results in Table 8. At the top quartile, the first stage is 32 percentage points, and the control school vaccination rate is 1 percentage points, so treatment raises school vaccination rates to 33 percentage points on average. At this high take-up rate, we find a small and insignificant impact on incidence. At other quartiles we continue not to find meaningful protective effects; at the bottom two quartiles we estimate large, noisy, and positive (i.e. damaging) effects. Thus the heterogeneity analysis does not support the hypothesis that the small school-level indirect effects are due to low school vaccination rates. An important limitation here is that even at the top quartile, the school vaccination rates are not especially high.

However, theoretical considerations also imply that the low school vaccination rate does not explain the null indirect effect. To show this, in Appendix E we simulate a standard susceptible-infected-recovered ("SIR") model, an epidemiological model of infection dynamics widely used by economists and others (e.g. Geoffard and Philipson (1997); May (2000); Jones et al. (2021); Farboodi et al. (2021); Goodkin-Gold et al. (2020, 2022), to investigate the consequences of a 20 percentage point increase in community vaccination rate. Under the assumptions of the SIR model—including that the vaccination prevents transmission and, critically, that people mix uniformly—we find that a 20 percentage point increase in community

²²Treatment group students go to middle schools, so we stratify the treatment group based on their school's vaccination rate. Control students go to elementary schools, with low vaccination rates throughout the sample period. To make treatment and control comparable within quartile, we stratify the control group based on the vaccination rate of the middle school they *would* go to if they were in seventh grade.

vaccination produces substantial indirect effects, except at very high levels of infectiousness.²³ However these high vaccination rates also imply that nearly all children eventually become infected, which we do not observe in our sample period.²⁴ As we also observe that vaccines reduce transmission at the household level, we conclude that differences in mixing—likely greater within household than within school—are an important factor in explaining the difference in indirect effects across our contexts.

6 Grade-wide Indirect Effects

So far we have shown that the protective benefits of vaccines spillover to other members of the same household but not to students in different grades in the same school. One explanation for the apparent absence of indirect effect in schools is that students in different grades do not come into close contact enough to make spillover protection important. We therefore turn to investigating within-grade spillovers. Specifically we examine the protective effects of belonging to a school *and grade* with a higher vaccination rate. To identify within-grade indirect effects, we take advantage of a regression discontinuity design created by birth date school entry cut-offs. Specifically, we compare COVID incidence between 7th graders (with mostly vaccine eligible peers) and 6th graders (with mostly vaccine ineligible peers), on either side of the date of birth cut-off that discontinuously shifts the probability of enrollment in grade 7 vs grade 6.

6.1 Empirical Strategy – Within grade indirect effect

We use a regression discontinuity (RD) design to measure possible indirect effects of the vaccine within grades by exploiting the fuzzy discontinuity in grade level assignment across date of birth. Specifically, in Indiana, school districts are required to offer Kindergarten to children who turn 5 on or before August 1

²³As Goodkin-Gold et al. (2020) explain, at very high levels of infectiousness, moderate vaccination produces no indirect effect because, although a vaccinated person will not herself infect an unvaccinated, the unvaccinated person is very likely to be infected by someone else.

²⁴Other evidence is consistent with relatively low infection and transmission in schools. In a contact tracing study of transmission among 191 students and school employees, Falk et al. (2021) find 5 percent of student cases attributable to in-school transmission, and no cases among school employees; Boutzoukas et al. (2022) also find low rates of in-school transmission. Likewise, Rosenberg et al. (2022) find in-person instruction is not associated with elevated COVID risk among college students in the Fall of 2020.

of the school year, and districts are allowed to enroll students turning 5 by October 1.²⁵ If all districts strictly enforced the August 1 cutoff, then in the 2021-2022 academic year, the grade 7 cohort would include children born between 2 August 2008 and 1 August 2009, while the grade 6 cohort would include children born August 2, 2009 through August 1, 2010. Let G_{it}^j be a dummy set to 1 if student i is enrolled in grade j or higher in academic year t . If we focus our attention on the population of students born within one year of 1 August 2009, the school starting age rule implies that $G_{i,2021-2022}^7 = 1(DOB_i \leq 1 \text{ August } 2009)$, which means that grade assignment is a discontinuous function of the date of birth. This discontinuity in grade assignment implies a discontinuity in the share of grade mates who are vaccine eligible (and ultimately vaccinated) in Fall 2021: all seventh graders are eligible for the vaccine at the start of the school year; in contrast only about a quarter of sixth graders are eligible then.

We exploit this discontinuity to study the indirect effects of peer vaccination on COVID infections among children born close to 1 August 2009. In a study population of children born within one year of 1 August 2009 attending Indiana schools in 2021-2022, we let dob_i represent the student's date of birth centered at 1 August 2009. Then the intent to treat (ITT) effect of enrollment in grade 7 (higher peer vaccination) rather than grade 6 (lower peer vaccination) is $\beta_{ITT} = \lim_{e \rightarrow 0} E[Y_i | dob_i = -e] - \lim_{e \rightarrow 0} E[Y_i | dob_i = +e]$. We estimate β_{ITT} by fitting a local weighted linear regression with the following form:

$$Y_i = \alpha_0 + \alpha_1 dob_i + \alpha_2 (dob_i \times G_i^7) + \beta_{ITT} G_i^7 + u_i$$

In practice our regression discontinuity design is fuzzy, not sharp, because compliance with the grade assignment rule is not perfect. This makes β^{ITT} difficult to interpret. We provide evidence on the first stage using data from Indiana's Department of Education on grade enrollment by date of birth. Unfortunately these data cannot be linked to vaccine take-up. We therefore estimate the reduced form and first stage separately.

We note also that, while we think of this approach as measuring within-grade indirect effects, these indirect effects arise from mixing both inside and outside the classroom. In particular, students in higher

²⁵<https://www.in.gov/doe/students/indiana-academic-standards/early-learning/kindergarten>.

grades will be in classrooms with more vaccinated students. Their social encounters outside of the classroom are also more likely to be with vaccinated friends, because (we hypothesize) students are more likely to form friendships with children in their own grade than in other grades.

We estimate the regression discontinuity models starting with a sample of 133,128 students born within 1 year of August 1, 2009 and alive on January 1, 2020. In placebo tests we also consider students born within 1 year of August 1, 2008 and 1 year of August 1, 2010. Because there is no time component to our identification strategy, we measure vaccination as of December 31, 2021, and we measure our main outcome (COVID incidence) as the fraction of months in Fall, 2021 in which the student had a positive COVID case. We estimate the model using triangular weights and select the bandwidth using the MSE-optimal bandwidth proposed by Calonico et al. (2014), and report bias-robust standard errors. We follow the best practices for graphical inferences as suggested by Korting et al. (2023).

To measure the “first stage” discontinuity in grade enrollment, we use administrative data from Indiana Department of Education. These data record, among other things, annual school enrollment information with student date of birth and grade level. Unfortunately our data use agreements do not permit us to link these data back to the INPC data. The most recent release of these data cover the 2019-2020 school year, so we measure the “first stage” as the likelihood of being in the higher grade in this year.²⁶

6.2 Results for within grade indirect effects

We provide summary statistics and evidence on the validity of the regression discontinuity design in Table 9, where we report the estimated discontinuities in pre-determined variables – race-ethnicity, sex, risk score, and fall 2020 COVID incidence. While we find a slight imbalance in the female proportion, we find no discontinuity in pre-period COVID incidence, pointing to little confounding. A basic threat to the validity of the regression discontinuity design is that school cut-offs may affect child health outcomes and behaviors independent of peer vaccination rates. To investigate this possibility, we construct a placebo

²⁶One might worry that, because some students skip grades and others repeat grades, the first stage deteriorates over time, making prior years unsuitable for measuring the first stage in 2021-22. Reassuringly, we find almost identical first stages using the 2018-2019 school year as the 2019-2020, implying little deterioration.

analysis using different cohorts: children born within 1 year of August 1 2010, or August 1, 2008. For these cohorts there is no first stage discontinuity in grade-wide vaccination rates, and Appendix Table A.7 shows that there is also no discontinuity in Covid incidence for these cohorts. These placebo results support the core exclusion restriction of our regression discontinuity design: there is nothing special about school cut-off sub-populations that puts them at higher or lower risk of COVID infection.

We show the key regression discontinuity results graphically in Figure 4. The top panel shows the first stage: enrollment in the higher grade falls by about 20 percentage points for students born just after August 1, relative to those born just before.²⁷ There is also a sharp decline at $d = 30$, implying that September 1 is a salient cutoff for many parents or school districts. As the seventh grade vaccination rate is about 20 percentage points higher than the sixth grade vaccination rate, this implies a first stage of about 4 percentage points, and an increase in the peer vaccination rate of 20 percentage points for the complier students. We caution that because we cannot link the vaccination and grade assignment data, we do not actually know the peer vaccination rate of complier students. It is possible that students are more likely to comply in high- or low-vaccine take-up schools, resulting in a slightly higher or lower first stage. Despite the discontinuously higher *peer* vaccination rate at the date-of-birth threshold, the middle panel shows that there is no discontinuity in *own* vaccination rate. Thus the protective effect of higher grade attendance, if present, would be due to peer vaccination and not own vaccination, i.e. an indirect effect. However the final panel shows no discontinuity in own COVID rates. The point estimates in Table 10 show a small and statistically insignificant discontinuity: assignment to an older grade, with a higher vaccination rate, reduces COVID incidence in Fall 2021 by a statistically insignificant 0.01 percentage points. While the point estimate is small, the confidence interval is fairly wide, ruling out reductions in COVID incidence larger than about 0.2 percentage points. This statistical uncertainty makes it difficult to draw firm conclusions. However the results here are inconsistent with the view that greater mixing within grades (in classrooms or with peers in the same grades) than across grades produces very large protective effects of vaccines.

²⁷We construct these RD figures following the guidelines of Korting et al. (2023).

7 Discussion and Conclusions

Using unique, large scale microdata from Indiana, we estimate direct and indirect effects of the COVID-19 vaccine, using the delayed approval of the vaccine for 11 year-olds relative to 12 year-olds as a source of variation in own and peer vaccine eligibility. This six-month approval delay leads to large and persistent differences in vaccination rates between 12 and 11 year-olds. We show that the higher vaccination rate of 12 year-olds reduces COVID incidence among 12 year-olds, with an implied vaccine effectiveness of 80 percent. We also find that the protective effect of vaccination extends to household members. Comparing adults who live with 12 year-olds to adults who live with 11 year-olds, we find that one household member's vaccine eligibility reduces COVID incidence among the rest of the household. The household-level indirect effect is quantitatively large: our baseline instrumental variables estimate implies that for each 12 year-old COVID case prevented by that 12 year-old's vaccination, an additional 0.6 cases are prevented among adults in her household. Because complier households—where kids become vaccinated when eligible—have high overall vaccination rates, these estimates mainly reflect a reduction in breakthrough infections, with little impact on severe COVID. However they show the possibility of substantial household spillovers in general.

Our final analyses ask how far these spillovers extend. We compare sixth graders in middle schools—whose schoolmates are mostly vaccine eligible—to sixth graders in elementary schools, with few vaccine eligible schoolmates. Despite going to school with 25 percent vaccinated peers, relative to under 5 percent, we detect no protective effect for the sixth graders in middle school, overall and throughout the distribution of middle school vaccination rates. While this finding of no school-level spillover could be explained by limiting mixing across grades, we also estimate a null (though imprecise) effect of increasing the share of grade-mates who are vaccine eligible.

Thus our results show that indirect effects of vaccines are context-specific, clear and substantial at the household level but small and insignificant at the school level. We caution that these results are subject to several caveats. First, our data are specific to a time and place: Indiana during fall 2021, when Delta was the primary circulating variant. Second, our estimates, like all estimates of vaccine effectiveness, are

contingent on ongoing mitigation measures and other behavioral responses. This may help explain why school-wide vaccination rates have little impact on transmission. Although we find no moderating effect of school mask mandates, schools took steps to reduce transmission by encouraging social distancing. Third, our estimates of vaccines' direct and indirect effects are derived from instrumental variables models and therefore reflect complier average causal effects. For estimating direct effects, the compliers are adolescents who become vaccinated soon after eligibility, and for indirect effects, they are households where adolescent COVID vaccination responds to eligibility. For other unvaccinated populations, treatment effects may be different. We hypothesize that our household indirect effects may substantially understate the indirect effects for the average household, because our compliers are likely to be vaccinated themselves, and we expect indirect effects to be more important among unvaccinated people. Finally, our ability to assign individuals to households is imperfect, likely producing measurement error in some key measures. Our investigation suggests the consequences of this measurement error is small, but it may another reason that lead us to understate household-level indirect effects.

A key question raised by our results is why indirect effects are substantial in some context but not others. In principle, the difference between the household and school-level indirect effect could reflect several factors. Standard SIR models highlight that the indirect effects of a vaccine depend on the vaccination rate (e.g. Goodkin-Gold et al. (2020)), and this differs substantially between our complier households and our treated middle schools. Another potential factor is differential mixing, which is not a factor in the textbook SIR model but is recognized by epidemiologists (Adam, 2020). We expect that household members interact much more closely on average than do schoolmates.

Our results point to mixing rather than vaccination rates as a key determinant of the indirect effect, for two reasons. First, we find no school-level indirect effect even for the highest vaccination rate schools. Second, in SIR models where a high vaccination rate is needed to obtain a large externality, low vaccination rates yield high rates of infection (Goodkin-Gold et al., 2020). But in practice transmission in schools appears low. The importance of household mixing echoes the recent finding of Ellen et al. (2023) that household crowding, more than census block density, is most associated with COVID risk.

The finding that indirect effects are larger for populations that mix more has implications for vaccine policy. The standard economic view is that vaccines should be subsidized or even mandated because vaccinations produce positive externality. The household-level indirect effect is consistent with that view. However, it may also be relatively straightforward for close contacts to privately negotiate optimal vaccination levels, following the logic of Coase (1960). Private-party solutions include families internalizing the family-level externalities, and employers using mandates or financial incentives to encourage vaccination. Our results also suggest, therefore, that efforts to ban employer mandates (especially in close-quarter settings) may be harmful. Whether families in fact internalize the externalities from vaccination, whether spillovers occur in other close settings of high policy importance such as healthcare, and whether organization/firm vaccination policies are an effective tool for internalizing them, are fruitful areas for future research.

References

- Abaluck, J., L. H. Kwong, A. Styczynski, A. Haque, M. A. Kabir, E. Bates-Jefferys, E. Crawford, J. Benjamin-Chung, S. Raihan, S. Rahman, et al. (2021). Impact of community masking on covid-19: A cluster-randomized trial in bangladesh. *Science* 375(6577), eabi9069.
- Acton, R. K., W. Cao, E. E. Cook, S. A. Imberman, and M. F. Lovenheim (2022). The effect of vaccine mandates on disease spread: Evidence from college covid-19 mandates. Technical report, National Bureau of Economic Research.
- Adam, D. (2020). Special report: The simulations driving the world's response to covid-19. *Nature* 580(7802), 316–319.
- Alexander, D. and E. Karger (2021). Do stay-at-home orders cause people to stay at home? effects of stay-at-home orders on consumer behavior. *The Review of Economics and Statistics*, 1–25.
- Allcott, H. (2015). Site selection bias in program evaluation. *The Quarterly Journal of Economics* 130(3), 1117–1165.
- Andersen, M. S., A. I. Bento, A. Basu, C. R. Marsicano, and K. I. Simon (2022). College openings in the united states increased mobility and covid-19 incidence. *medRxiv*, 2020–09.
- Angrist, J. D. and G. W. Imbens (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American statistical Association* 90(430), 431–442.
- Angrist, J. D., G. W. Imbens, and D. B. Rubin (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association* 91(434), 444–455.
- Baden, L. R., H. M. El Sahly, B. Essink, K. Kotloff, S. Frey, R. Novak, D. Diemert, S. A. Spector, N. Roupheal, C. B. Creech, et al. (2020). Efficacy and safety of the mrna-1273 sars-cov-2 vaccine. *New England journal of medicine*.
- Baker, J. M., R. M. Dahl, J. Cubilo, U. D. Parashar, and B. A. Lopman (2019). Effects of the rotavirus vaccine program across age groups in the united states: analysis of national claims data, 2001–2016. *BMC infectious diseases* 19(1), 1–11.
- Berry, C. R., A. Fowler, T. Glazer, S. Handel-Meyer, and A. MacMillen (2021). Evaluating the effects of shelter-in-place policies during the covid-19 pandemic. *Proceedings of the National Academy of Sciences* 118(15), e2019706118.
- Bleier, B. S., M. Ramanathan Jr, and A. P. Lane (2021). Covid-19 vaccines may not prevent nasal sars-cov-2 infection and asymptomatic transmission. *Otolaryngology–Head and Neck Surgery* 164(2), 305–307.
- Boutzoukas, A. E., K. O. Zimmerman, D. K. Benjamin, G. P. DeMuri, I. C. Kalu, M. J. Smith, K. A. McGann, S. Koval, M. A. Brookhart, and S. M. Butteris (2022). Secondary transmission of covid-19 in k–12 schools: findings from 2 states. *Pediatrics* 149(Supplement 2).

- Caddy, S. L. (2021). Coronavirus: Few vaccines prevent infection - here's why that's not a problem. *The Conversation*. January 5, 2021. Available at <https://theconversation.com/coronavirus-few-vaccines-prevent-infection-heres-why-thats-not-a-problem-152204>.
- Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica* 82(6), 2295–2326.
- Carmen, W., A. Elder, L. Wallace, K. McAulay, A. Walker, G. Murray, and D. Stott (2000). Effects of influenza vaccination of health-care workers on mortality of elderly people in long-term care: a randomized controlled trial. *Lancet* 355(9198), 93–7.
- Carpenter, C. S. and E. C. Lawler (2019). Direct and spillover effects of middle school vaccination requirements. *American Economic Journal: Economic Policy* 11(1), 95–125.
- Center for Disease Control and Prevention (2023). Underlying medical conditions associated with higher risk for severe covid-19: Information for healthcare professionals. <https://www.cdc.gov/coronavirus/2019-ncov/hcp/clinical-care/underlyingconditions.html>. Last Accessed March 20, 2023.
- Chan, T. Y. and B. H. Hamilton (2006). Learning, private information, and the economic evaluation of randomized experiments. *Journal of Political Economy* 114(6), 997–1040.
- Chan, T. Y., B. H. Hamilton, and N. W. Papageorge (2016). Health, risky behaviour and the value of medical innovation for infectious disease. *The Review of Economic Studies* 83(4), 1465–1510.
- Chetty, R., J. N. Friedman, N. Hendren, M. Stepner, and T. O. I. Team (2020). *How did COVID-19 and stabilization policies affect spending and employment? A new real-time economic tracker based on private sector data*, Volume 27431. National Bureau of Economic Research Cambridge, MA.
- Coase, R. H. (1960). The problem of social cost'(1960). *Journal of Law and Economics* 3, 1.
- COVID-19 School Data Hub (2022). CSDH Data. <https://www.covidschooldatahub.com>.
- Creech, C. B. (2022). Phone call with Daniel W. Sacks, August 5, 2022.
- Cronin, C. J. and W. N. Evans (2021). Total shutdowns, targeted restrictions, or individual responsibility: How to promote social distancing in the covid-19 era? *Journal of Health Economics* 79, 102497.
- Dave, D., A. I. Friedson, K. Matsuzawa, and J. J. Sabia (2021). When do shelter-in-place orders fight covid-19 best? policy heterogeneity across states and adoption time. *Economic inquiry* 59(1), 29–52.
- Ellen, I. G., R. Howland, and S. Glied (2023). Demons of density do higher-density environments put people at greater risk of contagious disease? *Journal of Housing Economics* 59, 101905.
- Falk, A., A. Benda, P. Falk, S. Steffen, Z. Wallace, and T. B. Høeg (2021). Covid-19 cases and transmission in 17 k–12 schools—wood county, wisconsin, august 31–november 29, 2020. *Morbidity and Mortality Weekly Report* 70(4), 136.

- Farboodi, M., G. Jarosch, and R. Shimer (2021). Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory* 196, 105293.
- Friedson, A. I., D. McNichols, J. J. Sabia, and D. Dave (2021). Shelter-in-place orders and public health: evidence from california during the covid-19 pandemic. *Journal of Policy Analysis and Management* 40(1), 258–283.
- Geoffard, P.-Y. and T. Philipson (1997). Disease eradication: private versus public vaccination. *The American Economic Review* 87(1), 222–230.
- Gillen, B., E. Snowberg, and L. Yariv (2019). Experimenting with measurement error: Techniques with applications to the caltech cohort study. *Journal of Political Economy* 127(4), 1826–1863.
- Ginther, D. K. and C. Zambrana (2021). Association of mask mandates and covid-19 case rates, hospitalizations, and deaths in kansas. *JAMA Network Open* 4(6), e2114514–e2114514.
- Goodkin-Gold, M., M. Kremer, C. M. Snyder, and H. Williams (2022). Optimal vaccine subsidies for endemic diseases. *International Journal of Industrial Organization* 84, 102840.
- Goodkin-Gold, M., M. Kremer, C. M. Snyder, and H. L. Williams (2020). Optimal vaccine subsidies for endemic and epidemic diseases. Technical report, National Bureau of Economic Research.
- Goolsbee, A. and C. Syverson (2021). Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. *Journal of public economics* 193, 104311.
- Gruber, J. (2005). *Public finance and public policy*. Macmillan.
- Gupta, S., T. Nguyen, S. Raman, B. Lee, F. Lozano-Rojas, A. Bento, K. Simon, and C. Wing (2021). Tracking public and private responses to the covid-19 epidemic: evidence from state and local government actions. *American Journal of Health Economics* 7(4), 361–404.
- Hall, W. D. (1999). Representation of blacks, women, and the very elderly (aged ≥ 80) in 28 major randomized clinical trials. *Ethnicity & disease* 9(3), 333–340.
- Halloran, M. E., M. P. Préziosi, and H. Chu (2003). Estimating vaccine efficacy from secondary attack rates. *Journal of the American Statistical Association* 98(461), 38–46.
- Humlum, M. K., M. Morthorst, and P. Thingholm (2022). Sibling spillovers and the choice to get vaccinated: Evidence from a regression discontinuity design.
- Imbens, G. W. and J. D. Angrist (1994). Identification and estimation of local average treatment effects. *Econometrica* 62(2), 467–475.
- Jones, C., T. Philippon, and V. Venkateswaran (2021). Optimal mitigation policies in a pandemic: Social distancing and working from home. *The Review of Financial Studies* 34(11), 5188–5223.
- Kates, J., C. Cox, and J. Michaud (2023). How much could covid-19 vaccines cost the u.s. after commercialization? Technical report, Kaiser Family Foundation.

- Kong, E. and D. Prinz (2020). Disentangling policy effects using proxy data: Which shutdown policies affected unemployment during the covid-19 pandemic? *Journal of Public Economics* 189, 104257.
- Korting, C., C. Lieberman, J. Matsudaira, Z. Pei, and Y. Shen (2023, 03). Visual Inference and Graphical Representation in Regression Discontinuity Designs*. *The Quarterly Journal of Economics*. qjad011.
- Leonhardt, D. (2021). Where are the tests? *New York Times*. The Morning Newsletter. September, 21, 2021.
- Lynge, F. P., K. Mølbak, M. Denwood, L. E. Christiansen, C. H. Møller, M. Rasmussen, A. S. Cohen, M. Stegger, J. Fonager, R. N. Sieber, et al. (2022). Effect of vaccination on household transmission of sars-cov-2 delta variant of concern. *Nature Communications* 13(1), 3764.
- Malani, A. (2008). Patient enrollment in medical trials: Selection bias in a randomized experiment. *Journal of Econometrics* 144(2), 341–351.
- Manski, C. F. and F. Molinari (2021). Estimating the covid-19 infection rate: Anatomy of an inference problem. *Journal of Econometrics* 220(1), 181–192.
- May, R. (2000). Simple rules with complex dynamics. *Science* 287(5453), 601–602.
- Mlcochova, P., S. A. Kemp, M. S. Dhar, G. Papa, B. Meng, I. A. Ferreira, R. Datir, D. A. Collier, A. Albecka, S. Singh, et al. (2021). Sars-cov-2 b. 1.617. 2 delta variant replication and immune evasion. *Nature* 599(7883), 114–119.
- National Center for Education Statistics (2022a). ELSI - elementary/secondary information system. <https://nces.ed.gov/ccd/elsi/>.
- National Center for Education Statistics (2022b). SABS - school attendance boundary survey. <https://nces.ed.gov/ccd/sabs/>.
- Peltzman, S. (1975). The effects of automobile safety regulation. *Journal of political Economy* 83(4), 677–725.
- Polack, F. P., S. J. Thomas, N. Kitchin, J. Absalon, A. Gurtman, S. Lockhart, J. L. Perez, G. P. Marc, E. D. Moreira, C. Zerbini, et al. (2020). Safety and efficacy of the bnt162b2 mrna covid-19 vaccine. *New England Journal of Medicine*.
- Regenstrief Institute (2022). RDS Data. <https://www.regenstrief.org/rds/data/>.
- Rosenberg, M., C. Chen, L. Golzarri-Arroyo, A. Carroll, N. Menachemi, and C. Ludema (2022). Sars-cov-2 reinfections in a us university setting, fall 2020 to spring 2021. *BMC Infectious Diseases* 22(1), 1–7.
- Sacarny, A., K. Baicker, and A. Finkelstein (2022). Out of the woodwork: enrollment spillovers in the oregon health insurance experiment. *American Economic Journal: Economic Policy* 14(3), 273–295.
- Sacks, D. W., N. Menachemi, P. Embi, and C. Wing (2022). What can we learn about sars-cov-2 prevalence from testing and hospital data? *The Review of Economics and Statistics*, forthcoming.

- Sadoff, J., G. Gray, A. Vandebosch, V. Cárdenas, G. Shukarev, B. Grinsztejn, P. A. Goepfert, C. Truyers, H. Fennema, B. Spiessens, et al. (2021). Safety and efficacy of single-dose ad26. cov2. s vaccine against covid-19. *New England Journal of Medicine* 384(23), 2187–2201.
- VanderWeele, T. J., E. J. T. Tchetgen, and M. E. Halloran (2012). Components of the indirect effect in vaccine trials: identification of contagion and infectiousness effects. *Epidemiology (Cambridge, Mass.)* 23(5), 751.
- Waldron, D. (2022a). Email correspondence with Kosali Simon, August 1, 2022.
- Waldron, D. (2022b). Covid-19 indiana schools. <https://github.com/dawaldron/covid-19-indiana-schools>, last accessed July 19, 2022.
- Ward, C. J. (2014). Influenza vaccination campaigns: is an ounce of prevention worth a pound of cure? *American Economic Journal: Applied Economics* 6(1), 38–72.
- Werner, J. M., A. Abdalla, N. Gara, M. G. Ghany, and B. Rehermann (2013). The hepatitis b vaccine protects re-exposed health care workers, but does not provide sterilizing immunity. *Gastroenterology* 145(5), 1026–1034.
- White, C. (2021). Measuring social and externality benefits of influenza vaccination. *Journal of Human Resources* 56(3), 749–785.

Table 1: Direct effects: Summary statistics and balance

Variable	Early Eligibility Group	Later Eligibility Group	Difference
Age	12.65 (0.48)	11.16 (0.37)	1.487 [0.002]
Female	0.49 (0.50)	0.48 (0.50)	0.008 [0.003]
White	0.69 (0.46)	0.70 (0.46)	-0.001 [0.003]
Black	0.14 (0.35)	0.14 (0.35)	0.001 [0.002]
Hispanic	0.10 (0.30)	0.10 (0.29)	0.002 [0.002]
Risk score	0.0014 (0.0005)	0.0014 (0.0005)	-0.0000 [0.0000]
Monthly COVID rate, pre-vaccine availability	0.0057 (0.0338)	0.0049 (0.0314)	0.0008 [0.0002]
# People	67,408	65,605	

Notes: Table reports the mean of the indicated variables, for the treatment and control group, as well as the difference in means. Each observation here is a person. Standard deviations in parentheses, heteroscedasticity robust standard errors in brackets. Pre-period COVID is the average monthly COVID incidence between June and December, 2020. Treatment group is born in the year ending May 12, 2009; control group is born in the year beginning November 3, 2009.

Table 2: Direct effects: Main estimates

Dep. var.	1+ dose	2+ doses	Any COVID	Non-COVID ER
DID estimate	0.2260 (0.0015)	0.2122 (0.0015)	-0.0027 (0.0002)	0.0004 (0.0003)
IV estimate			-0.0125 (0.0011)	0.0020 (0.0014)
Vaccine effectiveness			0.797 (0.018)	-0.164 (0.132)
# People				133,013

Notes: Table reports DID estimates of the effect of vaccine eligibility for each outcome, and DID-IV estimates and vaccine effectiveness for effect of two vaccine doses on monthly COVID incidence and ER visits. Treatment group is born in the year ending May 12, 2009; control group is born in the year beginning November 3, 2009. The post period is June-December, 2021, and the pre-period is June-December, 2020. Robust standard errors, clustered on individual, in parentheses. See Appendix B for details on vaccine effectiveness.

Table 3: Address-level indirect effects: Summary statistics and balance

Sample	Child with early eligibility present (1)	Child with early eligibility not present (2)	Difference (3)
Age	46.84 (12.18)	46.84 (12.18)	0.000 [0.079]
Female	0.58 (0.49)	0.57 (0.49)	0.010 [0.002]
White	0.77 (0.42)	0.78 (0.41)	-0.006 [0.003]
Black	0.12 (0.32)	0.12 (0.32)	0.004 [0.003]
Hispanic	0.08 (0.27)	0.07 (0.26)	0.005 [0.002]
Risk score	0.0197 (0.0189)	0.0197 (0.0187)	0.0000 [0.0001]
Monthly COVID rate, pre-vaccine availability	0.0172 (0.0579)	0.0171 (0.0581)	0.0001 [0.0004]
# People in household	4.89 (1.61)	4.92 (1.61)	-0.029 [0.015]
# People aged 14-18 in household	0.44 (0.65)	0.43 (0.66)	0.013 [0.006]
# People	65,521	59,113	
# Addresses	32,817	29,721	

Notes: Table reports the mean of the indicated variables for the indicated samples and the difference in means between the treatment group (column 1) and weighted control group (column 2). The treatment group consists of people aged 30 and older living with children born in the year ending May 12, 2009, and no children born in the next 18 months. Control group consists of people aged 30 and older living with children born in the year beginning November 3, 2009, and no children born in the prior 18 months. The control group is reweighted to match the age distribution in the treatment group. Pre-period COVID rate is measured August-December, 2020. Standard deviations in parentheses, and robust standard errors, clustered on household, in brackets.

Table 4: Address-level indirect effects: Main estimates

Outcome	Household Vaccination (1)	Own Vaccination (2)	Any COVID (3)	Severe COVID (4)
DID Estimate (Child eligibility)	0.10940 (0.00214)	0.00388 (0.00283)	-0.00089 (0.00034)	0.00011 (0.00014)
IV Estimate (Family Vaccination)			-0.00738 (0.00315)	0.00118 (0.00132)
N Addresses	62,538	62,538	62,538	62,538
N People	124,634	124,634	124,634	124,634

Notes: Each cell is a separate regression. The treatment group consists of people aged 30 and older living with children born in the year ending May 12, 2009, and no children born in the next 18 months. Control group consists of people aged 30 and older living with children born in the year beginning November 3, 2009, and no children born in the prior 18 months. The post-period is June-December, 2021, when older children become vaccine eligible. The pre-period is March, 2020 - May, 2021. For the DID rows, table reports the DID estimate of adolescent vaccine eligibility. For the IV rows, table reports the effect of household vaccination rate (excluding own vaccination) on the indicated outcomes. The control group is reweighted to match the age distribution of the treatment group. Robust standard errors, clustered on household, in parentheses.

Table 5: Address-level indirect effects: Heterogeneity

	First Stage (1)	IV- COVID (2)	IV- Severe COVID (3)	Vaccination Rate (4)	# People (5)	# Households (6)
<u>A. Split by risk score</u>						
Below median	0.1028 (0.0024)	-0.0135 (0.0044)	-0.0012 (0.0017)	0.46	69,719	48,647
Above median	0.1144 (0.0029)	-0.0005 (0.0040)	0.0039 (0.0020)	0.55	54,914	35,497
<u>B. Split by age</u>						
Below 45	0.0947 (0.0024)	-0.0134 (0.0050)	0.0006 (0.0020)	0.44	68,246	45,601
Above 44	0.1239 (0.0030)	-0.0018 (0.0037)	0.0017 (0.0017)	0.58	56,387	34,752
<u>C. Split by vaccination status as of May, 2021</u>						
Vaccinated	0.1759 (0.0029)	-0.0099 (0.0025)	0.0002 (0.0008)	1.00	51,308	32,591
Unvaccinated	0.0613 (0.0021)	-0.0027 (0.0077)	0.0029 (0.0035)	0.16	73,321	44,975

Notes: Each row reports IV estimates of the effect of family vaccination rate on own COVID incidence or severe COVID, for the indicated subsample. The design and sample are described in the notes to Table 4.

Table 6: School-level indirect effects: Summary statistics and balance

Variable	Middle Schoolers	Elementary Schoolers	Difference
Age	11.94 (0.70)	11.94 (0.70)	-0.002 [0.005]
Female	0.49 (0.50)	0.49 (0.50)	-0.006 [0.005]
White	0.76 (0.43)	0.74 (0.44)	0.022 [0.048]
Black	0.10 (0.29)	0.13 (0.34)	-0.037 [0.037]
Hispanic	0.09 (0.28)	0.10 (0.30)	-0.014 [0.024]
Risk score	0.0014 (0.0004)	0.0014 (0.0005)	-0.0000 [0.0000]
Monthly COVID rate, pre-vaccine availability	0.0072 (0.0379)	0.0072 (0.0378)	0.0000 [0.0007]
# People	34,511	21,380	

Notes: Table reports, for each variable, the mean among middle school sixth graders (treatment) and elementary school sixth graders (control), and the difference in means. Standard deviations are in parentheses and heteroscedasticity robust standard errors, clustered on school district, are in brackets. Schoolmates of middle school sixth graders but not elementary school sixth graders are vaccine eligible at the start of the 2021 school year. The “pre-vaccine” period is August-December, 2020 and the post-period is August-December, 2021.

Table 7: School-level indirect effects: Main estimates

Outcome	School Vaccination (1)	Own Vaccination (2)	Any COVID (3)
DID Estimate (Middle school effect)	0.22505 (0.01619)	0.00279 (0.00669)	0.00123 (0.00131)
IV Estimate (School vaccination rate)			0.00546 (0.00588)
N Students	55,891	55,891	55,891
N Districts	184	184	184

Notes: Table reports difference-in-difference estimates for the effect of middle school (rather than elementary school) on the indicated outcome and, for COVID incidence, 2SLS estimates of the effect of school vaccination (using the $\text{post} \times \text{treat}$ as the excluded instrument). The sample consists of monthly observations of sixth grade students with reliable school assignment, in the indicated months. The pre-period is 2020 and the post-period 2021. The treatment group is students who go to school with older grades (who are vaccine eligible in May 2021 and later the post-period) and the control group students who go to school with younger students (who are not vaccine-eligible until November 2021). Robust standard errors, clustered on school district, in parentheses.

Table 8: School-level indirect effects: Heterogeneity by middle school vaccination rate

	First Stage (1)	Reduced Form (2)	IV- COVID (3)	Control school Vaccination Rate (4)	# Students (5)	# Districts (6)
First quartile	0.1041 (0.0076)	0.0027 (0.0015)	0.0258 (0.0137)	0.01	14,272	94
Second quartile	0.1697 (0.0073)	0.0023 (0.0021)	0.0138 (0.0124)	0.01	13,943	61
Third quartile	0.2109 (0.0121)	-0.0009 (0.0013)	-0.0040 (0.0060)	0.03	14,228	36
Fourth quartile	0.3166 (0.0269)	-0.0001 (0.0013)	-0.0003 (0.0040)	0.01	13,448	21

Notes: Each row reports DID and IV estimates, stratified on quartile of middle school vaccination rate. Columns (1) and (2) are DID models, and column (3) an IV model. The outcome in column (1) is school vaccination rate, and in columns (2) and (3) it is COVID incidence. The DID and IV model are described in the notes to Table 7; see notes there for detail. Column (4) reports control group school vaccination rate in October, 2021. Robust standard errors, clustered on school district, in parentheses.

Table 9: Grade-level indirect effects: Summary statistics and balance tests

Variable	Control mean (1)	Discontinuity (2)	Standard Error (3)	Bandwidth (4)
COVID (Fall 2020)	0.005	-0.0003	(0.0008)	91.6
Black	0.146	-0.0004	(0.0109)	98.4
White	0.697	-0.0079	(0.0136)	102.7
Hispanic	0.097	0.0048	(0.0074)	113.9
Female	0.489	0.0010	(0.0144)	112.7
Risk score (x100)	0.138	0.0006	(0.0013)	105.9


Notes: Table presents balance tests for the regression discontinuity design. The running variable is date of birth relative to August 1, 2009. The first column presents the mean of the indicated variable in the control group (in the indicated bandwidth). The second column presents the point estimate and the third its robust standard error (as in Calonico et al. (2014)). The final column reports the MSE-optimal bandwidth of Calonico et al. (2014).

Table 10: Grade-level indirect effects: Regression discontinuity estimates

Outcome	In higher grade (1)	Own vaccination rate (2)	COVID incidence (3)
RD Estimate	-0.2248 (0.0269)	0.0126 (0.0119)	0.0002 (0.0011)
Bandwidth	42.3	63.9	128.7
# Observations	21,639	24,025	47,511

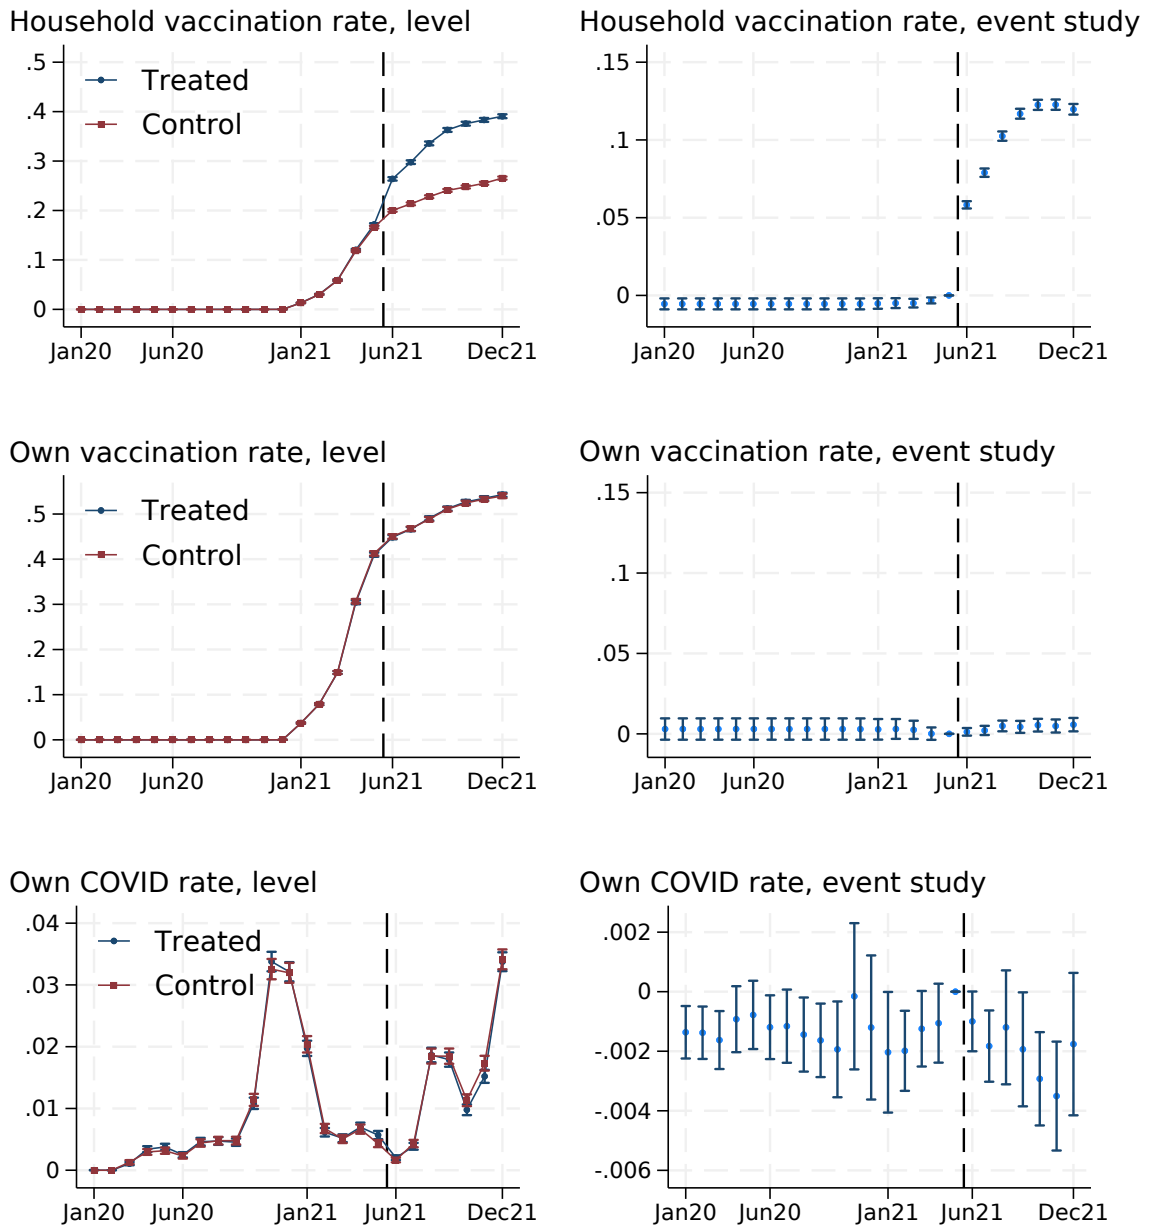
Notes: Table reports regression discontinuity estimates of the effect of being born before August 1, the cutoff for school starting age. Students with negative values of the running variable are more likely to be in 7th grade in the 2021-2022, and all 7th graders are vaccine eligible. The outcomes are as indicated; the vaccination rate is measured in December, 2021, and COVID incidence is averaged over August-December, 2021. The data underlying column (1) are administrative department of education records, which are only available through the 2019-20 school year. We therefore look at the probability that our cohort of students are in the higher grade (i.e. 5th grade instead of 4th for this year). The data in the other columns is from the INPC. The bandwidth is the Calonico et al. (2014) optimal one. We report the Calonico et al. (2014) robust standard errors, clustered on the running variable.

Figure 1: Direct effects: Event study



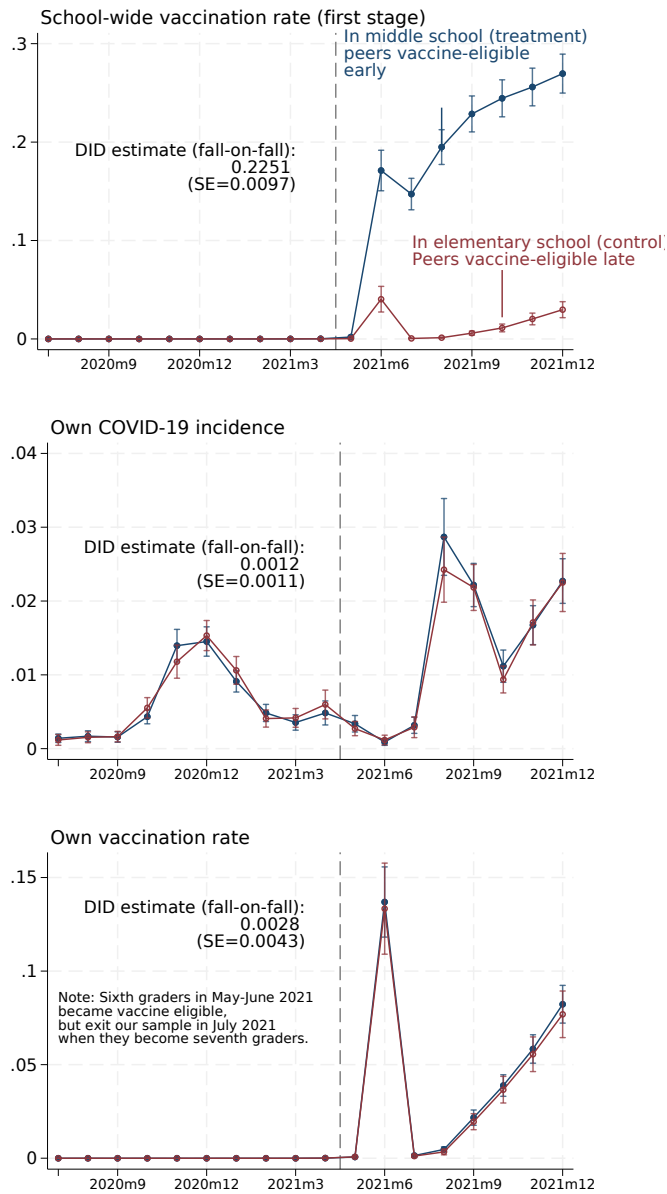
Notes: Figure plots means of the indicated variables, for Indiana residents born in the 12 months prior to May 12, 2009 (treatment) or the 12 months after November 3, 2009 (control), drawn from Regenstrief Institute data on Indiana COVID-19 vaccinations, COVID-19 testing, and emergency room visits. Bars indicate 95% confidence intervals, derived from robust standard errors clustered on individuals. (These confidence intervals are too small to be easily seen in the top panel.) The vertical line is the date the treatment group became vaccine-eligible. The DID estimate compares August-December 2020 and August-December 2021.

Figure 2: Household-level indirect effects: Event study




Notes: Figure plots the means of the indicated variables (on the left), and a visual event study (on the right) for the effect of having a vaccine eligible child in the household, among adults aged 30+. The treatment group consists of people aged 30+ living with kids turning 12 in the year leading up to May, 2021; these kids became eligible in May. The control group consists of people aged 30+ living with kids turning 12 in the year beginning November, 2021; they became eligible in November. We plot 95 percent confidence intervals, derived from standard errors clustered on household, as bars around the points. The family vaccination rate excludes own vaccinations.

Figure 3: School-level indirect effect: Event study



Notes: Figure plots means of the indicated variables, for Indiana sixth graders in the indicated school type. Sample is limited to students for whom we can reliably impute public school assignment. “Middle” schools are schools where the youngest grade is six, and “elementary” schools are where the oldest grade is six. Bars indicate 95% confidence intervals, derived from standard errors clustered on school district. The vertical line is the earliest date children were vaccine-eligible. The DID estimate compares August-December 2020 and August-December 2021.

Figure 4: Grade-level indirect effects: Regression discontinuity plots



Notes: Figure plots means of the indicated variables, in each bin of day of birth, relative to the cutoff birth date for school entry. The top panel shows the fraction of students in a higher grade (in this case, grade 6 in 2019-2020). Because higher-grade students are vaccine eligible in fall 2021, this panel shows an increase in the share of grade mates who are vaccine eligible. The middle panel shows the vaccination rate (as of December, 2021), and the final panel shows the rate of COVID incidence (fraction of months with any COVID) in Fall 2021.

A Appendix Exhibits

Table A.1: Sample restrictions and sample counts

Restriction	# People
<u>A. Direct effects design</u>	
Born May 13, 2008-May 12, 2009, or Nov 3, 2009-Nov 2, 2010	133,048
... and alive January 1, 2020	133,013
<u>B. Household-level indirect effects design</u>	
Age 30+, valid address, treated child or control child present	262,326
... and alive Jan-1-2020	262,326
... and just treated or just control child, and no interim children	246,106
... and 8 or fewer household members	165,024
<u>C. School-level indirect effects design</u>	
All 6th graders, 2020 or 2021	133,042
... and alive Jan-1-2020	133,005
... and has address	95,643
... and has school	91,182
... and school unambiguously middle or elementary	59,980
... and school not imputed	55,891
<u>D. Grade-level indirect effects design</u>	
Born within 1 year of August 1, 2009	133,165
... and alive Jan-1-2020	133,128

Notes: Table reports how the sample size changes as we impose our inclusion criteria for the different study designs. For the household-level indirect effects design, treated children are born May 13, 2008-May 12, 2009 and control children are born November 13, 2009 - Nov 2, 2010. The final row in each panel is the analysis sample.

Table A.2: Direct effects: Impact on ER visits with and without COVID

Type of ER visit	All	With positive test	With negative test	With no test (COVID-unrelated)
DID estimate	0.00040 (0.00031)	-0.00010 (0.00007)	-0.00014 (0.00013)	0.00042 (0.00029)
IV estimate	0.00188 (0.00145)	-0.00047 (0.00031)	-0.00064 (0.00061)	0.00199 (0.00139)
# People				133,013

Notes: Table shows the effect of vaccine eligibility (DID) and vaccine take-up (IV) on all ER visits, ER visits with positive COVID test in surrounding days, and ER visits with negative (and no positive) test in surrounding days, and ER visits with no COVID test in surrounding days, which we call “COVID-unrelated” visits. Surrounding days are 5 days before to four days after the ER visit. The sample and specification are defined in the notes to Table 2.

Table A.3: Direct effects: Sensitivity analysis

Dep. var.	1+ dose	2+ doses	Any COVID	Non-COVID ER
<u>A. Limit to born within 180 days of the cutoff</u>				
DID estimate	0.2207 (0.0022)	0.2090 (0.0021)	-0.0028 (0.0003)	0.0002 (0.0004)
IV estimate			-0.0134 (0.0016)	0.0010 (0.0020)
Vaccine effectiveness			0.818 (0.024)	-0.072 (0.161)
# People				64,420
<u>B. Expand to born within 540 months of cutoff</u>				
DID estimate	0.2290 (0.0013)	0.2149 (0.0012)	-0.0023 (0.0002)	0.0003 (0.0002)
IV estimate			-0.0108 (0.0009)	0.0012 (0.0009)
Vaccine effectiveness			0.757 (0.018)	-0.152 (0.130)
# People				195,730
<u>C. Has Address</u>				
DID estimate	0.2527 (0.0023)	0.2386 (0.0022)	-0.0030 (0.0003)	0.0006 (0.0004)
IV estimate			-0.0127 (0.0015)	0.0026 (0.0018)
Vaccine effectiveness			0.791 (0.025)	-0.215 (0.182)
# People				61,142

Notes: See notes to Table 2. Table reports identical estimates, except the sample differs as indicated. Specifically in panel A the sample is limited to children born within 180 days of the vaccine eligibility cutoff; in B it is extended to 540 days. In Panel C it is limited to students with an address (as used in the household-level indirect effects design).

Table A.4: Sources of household vaccination rate increases

Group	All ages (1)	Ages 11-13 (2)	Ages 14-18 (3)	Ages 0-10 (4)	Ages 18 and older (5)
DID Estimate (Child eligibility)	0.10940 (0.00214)	0.09785 (0.00085)	0.00446 (0.00078)	-0.00011 (0.00005)	0.00720 (0.00155)
Vaccination rate (Dec, 2021)	0.3475	0.0835	0.0444	0.0015	0.2180
N Addresses	62,538	62,538	62,538	62,538	62,538
N People	124,634	124,634	124,634	124,634	124,634

Notes: Table decomposes the impact of the presence of a vaccine eligible child on the household (leave-one-out) vaccination rate into vaccine responses by age group. The sample and specification are identical to those in Table 4. Each column reports the reduced form estimate for vaccination rates at a different age. Column (1) is the first stage of our IV model. Columns (2)-(5) aggregate to the first stage, as the denominator is always the household size minus 1, and the numerator always excludes the focal adult.

Table A.5: Household-level indirect effects: Sensitivity analysis

Outcome	Household Vaccination (1)	Own Vaccination (2)	Any COVID (3)	Severe COVID (4)
<u>A. Indirect effects among younger children</u>				
DID Estimate (Older sibling eligibility)	0.07498 (0.00304)	-0.00030 (0.00028)	-0.00113 (0.00046)	-0.00009 (0.00014)
IV Estimate (Family Vaccination)			-0.01778 (0.00648)	-0.00140 (0.00192)
N Addresses	28,578	28,578	28,578	28,578
N People	40,390	40,390	40,390	40,390
<u>B. Use 180 day window for child birthday</u>				
DID Estimate (Child eligibility)	0.10142 (0.00295)	0.00669 (0.00390)	-0.00111 (0.00047)	0.00017 (0.00020)
IV Estimate (Family Vaccination)			-0.01003 (0.00470)	0.00193 (0.00200)
N Addresses	32,706	32,706	32,706	32,706
N People	65,421	65,421	65,421	65,421
<u>C. Use 540 day window for child birthday</u>				
DID Estimate (Child eligibility)	0.11598 (0.00185)	-0.00024 (0.00244)	-0.00081 (0.00029)	0.00001 (0.00012)
IV Estimate (Family Vaccination)			-0.00704 (0.00252)	0.00007 (0.00105)
N Addresses	85,226	85,226	85,226	85,226
N People	169,630	169,630	169,630	169,630
<u>D. Obviously related instrumental variables</u>				
DID Estimate (Child eligibility)	0.12592 (0.00254)	0.00083 (0.00357)	-0.00097 (0.00048)	0.00005 (0.00019)
IV Estimate (Family Vaccination)			-0.00772 (0.00381)	0.00043 (0.00151)
N Addresses	51,936	51,936	51,936	51,936
N People	63,030	63,030	63,030	63,030

Notes: The sample and specification are identical to those in Table 4, except each panel differs in one dimension. In panel A the sample consists of children aged 0-10 (and alive on January 1, 2020), instead of adults 30 and older. In panel B we define treatment and control children as born within 180 days of the cutoff (vs. 365 at baseline); in panel C the definition is 540 days. In Panel D we exclude people whose address has any ambiguous matching, meaning we require that all residents of the address only have that address reported on the address log. In panel E we employ an obviously related instrumental variable strategy, using separate different time periods to measure the address for the instrument and the endogenous regressor.

Table A.6: School-level indirect effects: Robustness checks

Outcome	School Vaccination (1)	Own Vaccination (2)	Any COVID (3)
<u>A. Include students with imputed school assignment</u>			
DID Estimate (Middle school effect)	0.22427 (0.01556)	0.00252 (0.00641)	0.00122 (0.00122)
IV Estimate (School vaccination rate)			0.00546 (0.00550)
Test $\beta_1 = -0.0074$			0.019
N Students	59,240	59,240	59,240
N Districts	190	190	190
<u>B. Limit to students with unique school assignment</u>			
DID Estimate (Middle school effect)	0.22271 (0.01817)	0.00489 (0.00757)	0.00124 (0.00137)
IV Estimate (School vaccination rate)			0.00556 (0.00621)
Test $\beta_1 = -0.0074$			0.037
N Students	49,153	49,153	49,153
N Districts	176	176	176
<u>C. Control for instruction modality and mask mandate</u>			
DID Estimate (Middle school effect)	0.21868 (0.01688)	0.00106 (0.00721)	0.00152 (0.00125)
IV Estimate (School vaccination rate)			0.00694 (0.00579)
Test $\beta_1 = -0.0074$			0.013
N Students	52,968	52,968	52,968
N Districts	183	183	183


Notes: See notes to Table 7. The sample and specification are identical except as indicated. Panel A expands the sample to further include students with imputed school assignment. Panel B limits the sample to students with unique school assignment. Panel C indicators for the presence of virtual instruction and mask mandates. The sample changes because these are not observed for all schools. Each panel reports p-value of the test that the effect of school-wide vaccination rates, β_1 , equals -0.0081, the point estimate for the effect of household vaccination rates. Robust standard errors, clustered on school district, in parentheses.

Table A.7: Grade-level indirect effects: Placebo regression discontinuity designs

Outcome	Vaccination rate	COVID incidence
A. Born within 1 year of August 1, 2008		
RD Estimate	-0.0099 (0.0120)	0.0008 (0.0013)
Bandwidth	84.4	105.0
# Observations	32,104	39,601
B. Born within 1 year of August 1, 2010		
RD Estimate	0.0015 (0.0031)	-0.0014 (0.0013)
Bandwidth	66.2	100.3
# Observations	24,378	36,418


Notes: Table reports placebo tests for the effect of more vaccine-eligible grademates on the indicated outcome. The running variable is date of birth relative to the indicated date. These are placebo tests because there is no discontinuity in the share of grade mates that are vaccine eligible at these thresholds (in contrast to our main threshold of August 1, 2009). The vaccination rate is measured in December, 2021, and COVID incidence is averaged over August-December, 2021. The bandwidth is the Calonico et al. (2014) optimal one. We report the Calonico et al. (2014) robust standard errors, clustered on the running variable.

Figure A.1: COVID-19 vaccine eligibility by date of birth



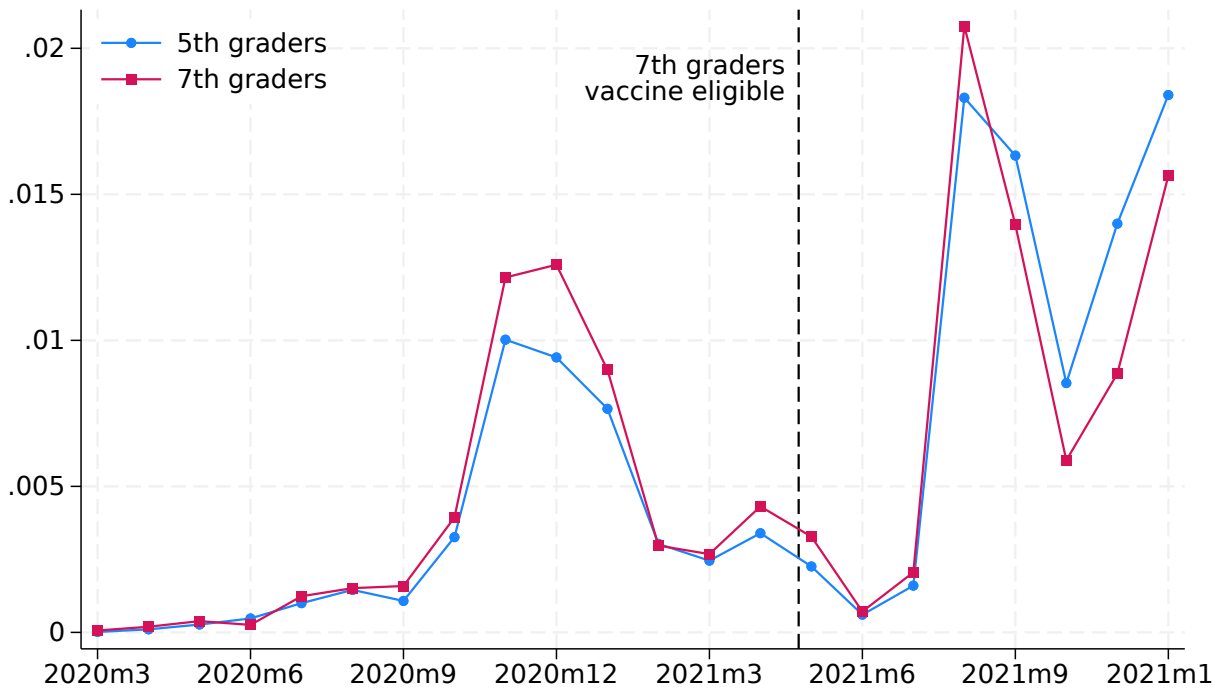
Notes: Figure shows the date (in 2021) Indiana expanded eligibility for the COVID-19 vaccines to different age groups, by day of birth.

Figure A.2: Complier Weights



Notes: Figure shows the distribution of weights underlying the interpretation of our 2SLS estimates as a weighted-average of complier marginal effects. The weight shows the weight put by the estimator on the fraction of compliers with a control ($z = 0$) vaccination rate below the indicated level, and a treatment ($z = 1$) vaccination rate at that level or greater. Our calculation of these weights follows Angrist and Imbens (1995).

Figure A.3: COVID incidence among fifth and seventh graders



Notes: Figure shows the fraction of 5th and 7th graders with at least one positive COVID test in each month.

B Vaccine effectiveness in an instrumental variables framework

Studies of causal effects in empirical microeconomics typically focus on treatment effect parameters that are expressed as differences in the expected value of treated and untreated potential outcomes for specified sub-populations. Instrumental variable estimators that account for incomplete take up or non-compliance with assigned treatments are interpreted as average causal effects among members of the complier sub-population. The clinical trials used to evaluate the effects of the COVID vaccines focused primarily on a somewhat different causal parameter, which is often referred to as “vaccine efficacy”.

In this appendix, we define a new parameter called “complier average vaccine efficacy” (CAVE). We derive an instrumental variables estimator of the CAVE that is valid under standard instrumental variable assumptions. We use the estimator in the paper to estimate the CAVE in our own-effects study design.

B.1 Notation and Assumptions

Use $i = 1 \dots N$ to index members of a study population. C_i is a binary observed outcome variable that indicates whether the person has a confirmed positive COVID test during a specified follow up window. V_i is a binary treatment variable indicating whether the person was vaccinated for COVID before the start of the follow up window. And Z_i is a binary instrumental variable, which is supposed to affect vaccine take up but is unrelated to COVID infection risk. Values of (C_i, V_i, Z_i) are observed for each member of the study population.

Observed vaccine take up and COVID infections are realizations of underlying potential outcomes. Specifically, let $V_i(z)$ be the vaccination status of person i when her instrument is set to z for $z = [0, 1]$. That means that realized vaccine take up is $V_i = V_i(0) + Z_i[V_i(1) - V_i(0)]$, where $V_i(1) - V_i(0)$ represents the causal effect of the instrument on person i 's vaccine take up. Similarly, let $C_i(z, v)$ be person i 's downstream COVID infection status if person i 's instrument is set to z and her vaccination status is set to v for $v = [0, 1]$.

We work with a set of five instrumental variable assumptions, which were originally described in papers by Imbens and Angrist (1994) and Angrist et al. (1996).

A1 SUTVA COVID infection outcomes are individualistic and do not depend on the vaccination status or instrumental variable assignments of any other members of the study population. More formally, let Z^{-i} be the $1 \times N - 1$ vector containing the instrumental variable assignments of each $j = 1 \dots N$ such that $j \neq i$. Likewise V^{-i} is the $1 \times N - 1$ vector of vaccination outcomes for each $j \neq i$. Now let $C_i(Z_i, V_i, Z^{-i}, V^{-i})$ be the potential outcome that person i would experience under a specific combination of own instrument and vaccine exposures **and** peer instrument and vaccine exposures. Under SUTVA $C_i(Z_i, V_i, Z^{-i}, V^{-i}) = C_i(Z_i, V_i)$ so that each person's potential outcomes do not depend on the vaccine status or instrumental variable status of any other member of the study population.

A2 Independence – The instrument is statistically independent of potential vaccine take up and potential COVID infection outcomes. Formally, independence implies $Pr(Z_i = 1 | V_i(z), C_i(z, v)) = Pr(Z_i = 1)$ for all combinations of z and v .

A3 Exclusion – The instrument has no causal effect on COVID infection outcomes. This implies that $C_i(z, v) = C_i(v)$ for all $i = 1 \dots N$.

A4 Monotonicity – The causal effect of the instrument on vaccine take up is non-negative for any individual in the sample. In other words $V_i(1) - V_i(0) \geq 0$ for all $i = 1 \dots N$.

A5 First Stage – The instrument has a non-zero causal effect on vaccine take up for at least some members of the study population so that $E[V_i(1) - V_i(0)] \neq 0$.

B.2 Treatment Effects

B.2.1 Additive Effects

At the person level, the additive causal effect of the vaccine on COVID infections is $\beta_i = C_i(1) - C_i(0)$. Since the infection variable is binary, the treatment effect for any single individual can only take on three different values. When $\beta_i = -1$, the person would have been infected with COVID if not for the vaccine. When $\beta_i = 1$ the person is infected with COVID if she is vaccinated but not infected if she is not vaccinated. Finally $\beta_i = 0$ if the person would either be infected in both vaccination states of the world or uninfected in both states of the world.

Treatment effect heterogeneity across subjects may occur for a variety of reasons, including: (i) behavioral responses to vaccination that lead some people to engage in riskier behaviors (Peltzman effects) or safer behaviors (health complementarity); (ii) biological differences in the immune response generated by the vaccine across subjects; and (iii) differences in epidemiological conditions (exposures) experienced by subjects in different times, places, and social settings.

The average treatment effect of the vaccine is

$$ATE = E[C_i(1) - C_i(0)].$$

The ATE is the difference in COVID infection rates between counterfactual states in which the population is universally vaccinated or universally unvaccinated. It is straightforward to define conditional average treatment effects. Standard examples are the average treatment effect on the treated: $ATT = E[C_i(1) - C_i(0) | V_i = 1]$, which represents the average effect of the vaccine on COVID infection among people who are actually vaccinated. If $ATT > ATE$, vaccinated people benefit more from the vaccine than unvaccinated people. If $ATT < ATE$ then vaccination would have larger effects on the unvaccinated population.

B.3 Vaccine Efficacy Effects

The literature on vaccine trials often focuses on measures of vaccine efficacy rather than on additive average treatment effects. Using the notation developed so far, vaccine efficacy is

$$\delta = 1 - \frac{Pr(C_i(1))}{Pr(C_i(0))}$$

With a vaccine that is perfectly effective, vaccinating the entire population eliminates 100% of the infections that would occur in the absence of the vaccine. Note, however, that vaccine efficiency is undefined when there is infection risk in the absence of the vaccine so that $Pr(C_i(0)) = 0$. In addition, it is less sensible to define efficacy at the person level the way we do for the additive treatment effect. For instance, $\delta_i = 1 - \frac{C_i(1)}{C_i(0)}$ will equal 0 for people who get infected regardless of vaccination status, 1 for

people who avoid an infection due to vaccination, and is undefined for people who are not infected in the absence of vaccination. That is unappealing since the vaccine could – in theory – increase infection risk among some people due to Peltzman type risk adjustment responses. The vaccine efficacy concept makes sense at a group level as long as there is a non-zero prevalence of cases of disease in the absence of vaccination.

B.4 Treatment Effects With Non-compliance

The COVID vaccine trials for the Pfizer, Moderna, and Johnson and Johnson vaccines used randomized experimental designs (Polack et al., 2020; Baden et al., 2020; Sadoff et al., 2021). People were randomly assigned to a vaccine group and a placebo group. COVID infections were measured at follow up and the infection rates in the two groups were used to estimate the causal effects of the vaccine. For example, Baden et al. (2020) report that at the end point of the Moderna trial, there were about 131.5 COVID cases per 10,000 people in the placebo group and about 7.8 COVID cases per 10,000 in the vaccine group. The average treatment effect implies that the vaccine reduced COVID infection rates by $7.8 - 131.5 \approx 123.7$ cases per 10,000. The efficacy of the vaccine was $1 - \frac{7.8}{123.7} \times 100 \approx 94.1\%$.

B.4.1 Complier Average Treatment Effects

The COVID vaccine trials experienced a small amount of non-compliance with the study protocol. Some subjects were lost to follow up, did not receive both doses of the vaccine, or experienced other events that made them ineligible. The main analysis in the trials used some form of per-protocol analysis in which these subjects were discarded, although various types of intent-to-treat samples were also considered.

In empirical economics, non-compliance with assigned treatments is often handled using instrumental variables analysis, providing a bridge between randomized experiments and quasi-experimental designs. A pair of papers by Imbens and Angrist (1994) and Angrist et al. (1996) show that in settings with a binary treatment and a binary instrumental variable satisfying assumptions A1-A5, the Wald-IV estimator identifies a parameter called the “Complier Average Treatment Effect” (CATE). Using the notation developed above, these papers show that

$$\frac{E[C_i|Z_i = 1] - E[C_i|Z_i = 0]}{E[Z_i|Z_i = 1] - E[Z_i|Z_i = 0]} = E[C_i(1) - C_i(0)|V_i(1) > V_i(0)]$$

The right hand side is the CATE, which is the average treatment effect in the sub-population of people who are induced to be vaccinated because of the instrumental variable. Given a valid instrumental variable, it is straightforward to estimate the CATE parameter from observed data. We report estimates of the CATE in our study of the own effects of the vaccine in Table 2.

B.4.2 Complier Vaccine Efficacy

In this section, we show how to identify a conditional version of the overall vaccine efficacy parameter, which we refer to as the “Complier Vaccine Efficacy” (CAE). The CAE is analogous to the CATE in the sense that it is a measure of vaccine efficacy in the sub-population of people who are induced to be vaccinated because of a binary instrumental variable. The CAE parameter that we focus on in this section is defined as:

$$\delta_{complier} = 1 - \frac{Pr(C_i(1)|V_i(1) > V_i(0))}{Pr(C_i(0)|V_i(1) > V_i(0))}.$$

$\delta_{complier}$ is a function of two counterfactual quantities. $Pr(C_i(0)|V_i(1) > V_i(0))$ is the complier base rate: it represents the COVIDinfection rate among compliers in the absence of vaccination. $Pr(C_i(1)|V_i(1) > V_i(0))$ is the complier breakthrough rate. It represents the complier infection rate when the compliers are vaccinated.

In this section, we show that both of these quantities are identified under assumptions A1-A5. The CAE is identified under the additional restriction that $Pr(C_i(0)|V_i(1) > V_i(0)) > 0$.

The First Stage

Under A1-A5, the first stage comparison identifies the fraction of compliers in the population:

$$\begin{aligned} F &= E[V_i|Z_i = 1] - E[V_i|Z_i = 0] \\ &= E[V_i(1)|Z_i = 1] - E[V_i(0)|Z_i = 0] \\ &= E[V_i(1)] - E[V_i(0)] \\ &= P[V_i(1) > V_i(0)] \end{aligned}$$

The second equality follows after substitution of the potential vaccine take up expression for the observed vaccine take up outcomes. The third equality imposes the independence assumption. And the fourth equality imposes the monotonicity condition. This shows that the first stage difference in vaccine take up rates identifies the prevalence of compliers.

The Complier Base Rate

The logical challenge in identifying the complier base rate is that complier status is unknown at the individual level, and unvaccinated COVID potential outcomes are not observed for the full population. We can apply the standard instrumental variables analysis to an adjusted/censored outcome variable to uncover complier averages of the individual outcomes.

Let $R_i^{base} = (1 - V_i)C_i$ to be an adjusted outcome that is set to 0 for people who are vaccinated and set to the value of C_i for people who are unvaccinated. The reduced form difference in (adjusted) COVID outcomes across levels of the instrument is:

$$\begin{aligned} ITT_{base} &= E[R_i^{base}|Z_i = 1] - E[R_i^{base}|Z_i = 0] \\ &= E[(1 - V_i)C_i|Z_i = 1] - E[(1 - V_i)C_i|Z_i = 0] \\ &= E[(1 - V_i(1))C_i(0)|Z_i = 1] - E[(1 - V_i(0))C_i(0)|Z_i = 0] \\ &= E[C_i(0)(V_i(0) - V_i(1))] \\ &= -E[C_i(0)|V_i(1) > V_i(0)]Pr(V_i(1) > V_i(0)). \end{aligned}$$

The second equality substitutes the definition of the adjusted outcome, and the third equality introduces the potential outcomes structure, invoking the exclusion restriction. The fourth equality imposes the independence assumption to drop conditioning on the instrument. The fifth line decomposes the ex-

pectation using the fact that $V_i(0) - V_i(1)$ can only take on the values 1, 0, and -1. Two of the three terms drop out: the zero term is multiplied by zero and $Pr(V_i(0) - V_i(1) = 1) = 0$ under Under A4 (monotonicity). Thus ITT_{base} is equal to the negative of the complier base rate multiplied by the prevalence of compliers. Dividing by the negative of the complier share using a standard Wald Ratio gives:

$$\begin{aligned}
W_{base} &= \frac{ITT_{base}}{-F} \\
&= \frac{E[R_i^{base}|Z_i = 1] - E[R_i^{base}|Z_i = 0]}{-(E[V_i|Z_i = 1] - E[V_i|Z_i = 0])} \\
&= \frac{-E[C_i(0)|V_i(1) > V_i(0)]Pr(V_i(1) > V_i(0))}{-Pr(V_i(1) > V_i(0))} \\
&= E[C_i(0)|V_i(1) > V_i(0)] \\
&= Pr[C_i(0) = 1|V_i(1) > V_i(0)].
\end{aligned}$$

The Complier Breakthrough Rate

Following a parallel approach for the complier breakthrough rate, define the adjusted outcome $R_i^{break} = V_i C_i$, which is set to 0 for people who are unvaccinated and set to C_i for people who are vaccinated. The reduced form comparison in this case is:

$$\begin{aligned}
ITT_{break} &= E[R_i^{break}|Z_i = 1] - E[R_i^{break}|Z_i = 0] \\
&= E[V_i C_i|Z_i = 1] - E[V_i C_i|Z_i = 0] \\
&= E[V_i(1)C_i(1)|Z_i = 1] - E[(V_i(0)C_i(1)|Z_i = 0] \\
&= E[C_i(1)(V_i(1) - V_i(0))] \\
&= E[C_i(1)|V_i(1) > V_i(0)]Pr(V_i(1) > V_i(0)).
\end{aligned}$$

Here, the second equality uses the definition of the adjusted outcome and the third line equality introduces the potential outcomes structure, invoking the SUTVA condition and the exclusion restriction. The fourth equality imposes the independence assumption and collects terms. The fifth line decomposes the expected value of the product of $C_i(1)$ and $V_i(1) - V_i(0)$ and imposes the monotonicity assumption. The result shows that ITT_{break} is the complier breakthrough infection rate multiplied by the prevalence of compliers. The Wald ratio isolates the complier breakthrough rate:

$$\begin{aligned}
W_{break} &= \frac{ITT_{break}}{F} \\
&= \frac{E[R_i^{break}|Z_i = 1] - E[R_i^{break}|Z_i = 0]}{E[V_i|Z_i = 1] - E[V_i|Z_i = 0]} \\
&= \frac{E[C_i(1)|V_i(1) > V_i(0)]Pr(V_i(1) > V_i(0))}{Pr(V_i(1) > V_i(0))} \\
&= E[C_i(1)|V_i(1) > V_i(0)] \\
&= Pr[C_i(1) = 1|V_i(1) > V_i(0)].
\end{aligned}$$

Identifying Complier Vaccine Efficacy

The complier average vaccine efficiency can be estimated using the ratio of the two Wald ratios:

$$\begin{aligned}
\delta_{complier} &= 1 - \frac{W_{break}}{W_{base}} \\
&= 1 - \frac{ITT_{break} \times F^{-1}}{-ITT_{base} \times F^{-1}} \\
&= 1 + \frac{ITT_{break}}{ITT_{base}} \\
&= 1 - \frac{Pr(C_i(1)|V_i(1) > V_i(0))}{Pr(C_i(0)|V_i(1) > V_i(0))}.
\end{aligned}$$

Interestingly, the first stages cancel and so the efficiency is equal to 1 plus the ratio of the reduced forms. In practice, you could estimate the complier efficiency by computing the two IV estimates (complier base rate and complier breakthrough rate) directly and then computing the ratio of the two. Or you could compute the two ITT effects and compute their ratio. In both cases, it would be sensible to use a stacked framework so that one could straightforwardly produce a joint covariance matrix.

C Selection into testing

Our primary outcome is an indicator for at least one lab confirmed case of COVID-19 in a given month, $Pr(\text{test positive})$. This decomposes as

$$Pr(\text{test positive}) = Pr(\text{positive}|\text{test} = 1) \cdot Pr(\text{test} = 1). \quad (7)$$

This decomposition shows that our outcome can change, in principle, not because of true changes in COVID incidence but because of changes in testing behavior, i.e. changes in $Pr(\text{test} = 1)$.

Here we argue that changing test behavior is unlikely to account for our key qualitative results. Key to our argument is the observation that the test positivity rate, $Pr(\text{positive}|\text{test} = 1)$, sometimes called the yield, reflects the combination of selection into testing and overall COVID incidence (Manski and Molinari, 2021; Sacks et al., 2022). Holding fixed COVID incidence, as the tested population becomes more positively selected, $Pr(\text{positive}|\text{test} = 1)$ increases. Thus our estimate that vaccine eligibility reduces measured COVID incidence might be explained by reduced testing rather than reduced incidence.

If this explanation were true, we would expect to see that vaccine eligibility increases the test positivity rate, because the marginal patient induced not to test by vaccine eligibility should have a relatively low chance of having COVID. More generally, strong selection effects imply opposite signed effects on the unconditional probability of a positive test and on the positivity rate. We therefore re-estimate our main DID models, for direct effects, indirect household-level effects, and indirect-school level effects, but looking at test positivity as an outcome. We define test positivity as the fraction of positive COVID tests in a given month, as a share of all COVID tests.

The results, in Appendix Table C.1, are inconsistent with strong selection effects. We report DID estimates in the first row and, to contextualize the magnitudes, we report main effects for “post” and “treat” in the remaining rows. In general treatment effects are small relative to main effects; positivity spiked during the delta wave. The results in column (1) show that own vaccine eligibility reduces test positivity. This is of course consistent with our conclusion that the vaccine is effective for the vaccinated. But it is inconsistent with falling test rates (conditional on symptoms) among the vaccinated, and thus indicates that changing selection into testing does not account for our finding of substantial direct effects. Column (2) shows a small and insignificant effect of household member eligibility on positivity, again consistent with no change in selection behavior. Column (3) shows a positive and marginally significant effect of schoolmate eligibility on positivity. Given that we find a positive but insignificant effect of schoolmate eligibility on own COVID incidence, this result too is inconsistent with important selection effects.

Table C.1: Impact of own and peer vaccine eligibility on test positivity

Design	Direct effect (1)	Indirect - household (2)	Indirect - school (3)
DID Estimate	-0.034 (0.007)	-0.003 (0.005)	0.030 (0.016)
Coef. on post	0.115 (0.005)	0.076 (0.004)	0.091 (0.014)
Coef. on treat	0.011 (0.004)	0.003 (0.003)	-0.005 (0.009)
Constant	0.110 (0.003)	0.147 (0.002)	0.119 (0.007)
# Observations	43,948	113,782	15,472
# Clusters	26,469	45,368	326
Clustering	Individual	Household	School

Notes: Table difference-in-difference estimates for the effect of own (column 1), household member (column 2), or schoolmate (column 3) vaccine eligibility on own test positivity rate. See notes to Tables 1, 3, and 6 for description on each sample and design. The sample here is further limited to person-months with at least one COVID test, for whom positivity is defined. Robust standard errors, clustered at the indicated level, in parentheses.

D Average Causal Response and Vaccine Effects

When we examined the direct effect of the vaccine treatment, a person’s treatment status was binary – vaccinated or unvaccinated. It made sense to think about a pair of potential health outcomes for the person under vaccinated or unvaccinated conditions. Indirect vaccine effects – at least in the contexts we consider in this paper – are better understood as a treatment with variable intensity. We want to understand the same person’s health outcomes under alternative vaccination rates among some specified peer group. In section 4 we estimate household indirect effects using two stage least squares regressions. The estimating equations we use are:

$$Y_{it} = \beta_0 + \beta_1 \widehat{Vacc}_{-it} + \beta_2 treat_i + \beta_3 post_t + \epsilon_{it}$$

$$Vacc_{-it} = \alpha_0 + \alpha_1 EarlyEligChildPresent_i \cdot post_t + \alpha_2 EarlyEligChildPresent_i + \alpha_3 post_t + \epsilon_{it}$$

In the model, $Vacc_{-it}$ is the vaccination rate of all members in i ’s household except i herself, and the excluded instrument is the interaction $EarlyEligChild_i \cdot post_t$. Our primary outcome— Y_{it} —is a dummy variable indicating that the person had a lab-confirmed COVID infection during month t . The first stage is a difference-in-difference regression that measures the effect of the early eligible child on the vaccination rate in the household where the treated adult is living.

On the surface, the two stage least squares framework appears to impose unrealistic linearity assumptions, which would imply that increases in household vaccination rates have a constant effect on COVID infection risk. It turns out, however, that the two stage least squares coefficient – β_1 – can be interpreted through the lens of the average causal response theorem developed in Angrist and Imbens (1995). In particular, in a model where the causal effects of marginal increases in vaccination rates are heterogeneous across individuals and also across levels of the vaccination rate, β_1 represents a weighted average of causal effects of increasing the household vaccination rate among complier households, averaging over the distribution of instrument-induced changes in household vaccination.

To see the idea, suppose that the household vaccination rate $Vacc_{-it}$ takes values on a grid of $j = 0 \dots 1000$ different values. The values on the grid correspond to household vaccination rates ranging from 0% vaccinated to 100% vaccinated in steps of 0.1 percentage points. Thus v_j is the household vaccination rate at grid point j , which means that $v_0 = 0, v_1 = .001, \dots v_{1000} = 1$.

Abstracting from covariates and letting $Z_i = treat_i \cdot post_t$ represent the value of the instrumental variable for person i . $V_{-it}(z)$ represents the potential household vaccination rate outcome that person i would experience when exposed to instrumental variable setting $Z_i = z$. For instance, $V_{-it}(1)$ represents the household vaccination rate the person would face given exposure to an early eligible child, and $V_{-it}(0)$ represents the same person’s household vaccination rate when exposed to late eligible child. Finally, let $Y_i(z, v_j)$ represent the health outcome that person i would experience if she were exposed to instrument setting $Z_i = z$ and household vaccination rate $V_{-it} = v_j$.

Impose the standard instrumental variable assumptions:

- I. SUTVA While our household indirect effects analysis is designed to measure spillover effects within households, we do maintain the assumption that the Covid-19 infection outcomes of adults in our household spillover sample are household specific and do not depend on the vaccination status or instrumental variable assignments of any other households in the study population. More for-

mally, let HZ^{-i} be the $1 \times N - 1$ vector of instrumental variable of households that do not contain person i . Likewise HV^{-i} is the $1 \times N - 1$ vector of vaccination outcomes of the adult members of households that do not contain person i . Now let $Y_i(Z_{-i}, V_i, HZ^{-i}, HV^{-i})$ be the potential outcome that person i would experience under a specific combination of own instrument and vaccine exposures **and** peer instrument and vaccine exposures. Under SUTVA $Y_i(Z_i, V_i, Z^{-i}, V^{-i}) = Y_i(Z_i, V_i)$ so that each person's potential outcomes do not depend on the vaccine or instrumental status in other households.

2. Independence – $Pr(Y_i(z, v_j), V_{-it}(z)|Z_i) = Pr(Y_i(z, v_j), V_{-it}(z))$ for all z, v_j
3. Exclusion – $Y_i(z, v_j) = Y_i(v_j)$
4. First Stage – $E[V_{-it}|Z_i = 1] \neq E[V_{-it}|Z_i = 0]$
5. Monotonicity – $V_{-it(1)} \geq V_{-it(0)}$ for all i

Under SUTVA and the exclusion restriction, $Y_i(v_j)$ represents the potential outcome that person i would experience if she were exposed to a household vaccination rate of $V_{acc-it} = v_j$. And $\tau_i^{j,j-1} = Y_i(v_j) - Y_i(v_{j-1})$ is causal effect of increasing the household vaccination rate from v_{j-1} to v_j on person i 's COVID related health outcome. Angrist and Imbens (1995) show that under the instrumental variable assumptions above, the conventional Wald ratio corresponds to a weighted average of these incremental causal effects among a collection of complier groups affected by the instrument. Specifically, they show that:

$$\frac{E[Y_i|Z_i = 1] - E[Y_i|Z_i = 0]}{E[V_{acc-i}|Z_i = 1] - E[V_{acc-i}|Z_i = 0]} = \sum_{j=1}^J w_j \times E[\tau_i^{j,j-1} | V_{-it}(1) \geq j > V_{-it}(0)].$$

Angrist and Imbens (1995) develop their result for Wald estimators whereas we estimate 2SLS models. However our 2SLS estimates are numerically identical to Wald estimates applied to first differenced data, taking the difference at the individual level between average outcomes in the pre and post period, and using treatment as an instrument for the difference in household vaccination rate. The expression above implies that the Wald Ratio gives a weighted average of the causal effects of incremental changes in household vaccination from across the range of possible vaccination rates. The weights in the expression— w_j —are proportional to the fraction of compliers whose household vaccination rate increases from below v_j to v_j or above. Using the notation developed above and abstracting from covariates, the weights are

$$w_j = \frac{Pr(V_{-it}(1) \geq j > V_{-it}(0))}{\sum_j^J Pr(V_{-it}(1) \geq j > V_{-it}(0))}$$

Figure A.2 plots the weights from our household indirect effects analysis. The figure shows that there is a positive weight on household vaccination rates across the full 0-1 range. However most of the weight is concentrated on household vaccination rates between 33% and 75%. Very little weight is applied to household vaccination rates below 25%. This distribution of weights implies that our two stage least squares estimate of the household indirect effect is a close approximation to the the average casual effect of marginal changes in household vaccination rates for intermediate vaccination rates. Our estimate may

be less informative about indirect effects created when a household shifts from a very low vaccination rate to a moderate moderate vaccination rate.

E A quantitative model of vaccine externalities

This section presents and numerically analyzes an SIR-model with vaccination to understand whether incomplete vaccine take-up could explain the near-zero spillovers we estimate. We have three results. First, except for very high levels of infectiousness, prevalence among the unvaccinated is nearly linear in the vaccination rate, up to the herd immunity threshold. The effect of additional vaccinations on the unvaccinated therefore is not very sensitive to baseline vaccination rates except at high levels of infectiousness. Our second result is that at high levels of infectiousness, a marginal vaccination provide little protection to the unvaccinated, because an unvaccinated person is likely to become infected from another source.²⁸

Taken together these results imply that when spillovers exist, they are likely large enough for us to detect from a 20 percentage point increase in schoolmate vaccination rates. Our third result shows this directly: at all levels of vaccination below herd immunity, the simulated effect of a 20 percentage point increase in vaccination is much larger than what our confidence intervals rule out, except in the case of high infectiousness, when spillovers are small.

We caution that this model is particularly simple and may not capture the dynamics of COVID-19. The results here do not necessarily generalize to other disease models.

E.1 Model set-up

We consider the simple SIR model with uniform mixing and a share v of the population of size N is vaccinated. The vaccine is assumed to be 100 percent effective, and we model vaccinated people as removed from the susceptible pool. Our set up is the discrete time analog of the model in Goodkin-Gold et al. (2020), except we assume perfect effectiveness and abstract from the vaccine demand phase. Thus

$$N = S + I + R + vN.$$

The equations describing infection dynamics are

$$\begin{aligned}\Delta S &= -\beta S \cdot I/N \\ \Delta I &= \beta S \cdot I/N - \gamma I \\ \Delta R &= \gamma I.\end{aligned}$$

Here β is the transmission rate and γ is the recovery rate.

A key parameter is the reproductive number \mathcal{R} , the number of new infections spawned by a single infection. The basic reproductive number \mathcal{R}_0 is the value of \mathcal{R} in a completely susceptible population, so $\mathcal{R}_0 = \beta/\gamma$. When $\mathcal{R} < 1$, infections do not replace themselves and so disease outbreaks die out.

Because the vaccinated and unvaccinated populations mix uniformly, a vaccination rate of v scales down the susceptible population by $(1 - v)$, and so reduce the reproductive number $(1 - v)$. A large

²⁸This intuition is from Goodkin-Gold et al. (2020), who develop it in a series of related results.

enough vaccinated population ensures that $R < 1$; the so-called herd-immunity vaccination rate guaranteeing this condition is

$$v^* = 1 - 1/R_0.$$

As we will see, this threshold plays an important role in the results.

Simulation details Closed-form solutions for final infection rates and infection dynamics do not exist, so we solve the model with forward simulation to obtain the final-period count of ever infected individuals, $R(T)$. We simulate for $T = 20000$ time periods, starting with $I(1) = 1$ and $R(1) = 0$. In each run we verify that the simulation converges in the sense that the number of infected people change by less than $1/1000$ over the last periods.

Parameterization The model parameters are N , β , γ and v . We fix $N = 100,00$ and $\gamma = 1/10$. We choose β so that $\mathcal{R}_0 \in 1.1, 1.5, 2, 3, 5$; note that $\gamma = .1$, meaning a 10 day expected infection length. For each β I vary the vaccination rate from 0 to 1 in increments of 0.01.

These parameters trace out a range of reasonable values for \mathcal{R}_0 in the context of COVID-19; 1.1 is lower than estimates; 1.5 is the estimated \mathcal{R}_0 for the ancestral strain (also used in Goodkin-Gold et al. (2020)), and 5 represents a very high estimate, possibly occurring with the latest strains, although it is unclear if high transmission of the latest waves reflects immune escape or high \mathcal{R}_0 . The scale of γ is not relevant for \mathcal{R}_0 , but $\gamma = 1$ has multiple advantages. First, high values of γ ensure that the epidemic concludes in relatively few iterations. However, β must be less than 1 since it is a transmission probability. Choosing $\gamma = .1$ means that $\beta = .5$ when $\mathcal{R}_0 = 5$.

Simulation output: For each value of v and \mathcal{R}_0 we calculate the fraction of the unvaccinated population that ever becomes. Since the vaccinated cannot be infected, this fraction is

$$pr(infected|unvaccinated; v, R_0) = R(T; v, R_0)/(N \cdot (1 - v)).$$

Our empirical analysis of vaccine spillovers considers a shock that increases peer vaccination rate by roughly 20 percentage points. We therefore also calculate the implied impact of such a shock on the unvaccinated infection rate:

$$\Delta pr(infected|unvacc; v, R_0) = pr(infected|unvacc; v + .2, R_0) - pr(infected|unvacc; v, R_0).$$

E.2 Model results


We begin by showing the fraction of the unvaccinated population that ever becomes infected, as a function of the vaccination rate, for various \mathcal{R}_0 , in Appendix Figure E.1. Several patterns are clear in the figure. Most obviously, marginal vaccinations beyond the herd immunity threshold (the vertical lines) have only very small impacts on the unvaccinated, because at the the herd immunity threshold and beyond, infections die out and nearly all unvaccinated would not become infected even absent greater vaccination.

More importantly, for low levels of infectiousness— $\mathcal{R}_0 < 3$ —the relationship between $Pr(infected|unvacc)$ and v is approximately linear, up to the herd immunity threshold. Thus the marginal benefit of vaccinations is roughly constant in v ; it does not depend on the starting level of vaccination. Estimates of the impact of greater peer vaccination on own infections, if this model were true, would not be too sensitive to baseline vaccination rate. For high levels of infection, the nonlinearity is stronger. However it is also true at these high levels of infection, especially when $\mathcal{R}_0 = 5$, there is very little external benefit of vaccines; even large increases in the vaccination rate do not produce large reductions in unvaccinated

incidence, except at very high levels of vaccination.


We show this more specifically in Appendix Figure E.2. The figure shows the simulated impact of a 20 percentage point increase in the vaccination rate on incidence among the unvaccinated, as a function of the initial vaccination rate, for different levels of \mathcal{R}_0 . While the effect size does vary with v , it is always large when $\mathcal{R}_0 < 5$. Indeed the lower bound of the confidence interval from our main estimates—about -0.005 — easily lets us rule out any effect size in the figure, except when either (a) herd immunity is reached, or (b) infectiousness is so high that the spillover is small for a wide range of initial vaccination levels. Even in the case, however, the implied effect is on the order of a few percentage points, an order of magnitude larger (in absolute value) than the lower bound of our confidence interval.

Figure E.1: Infections among the unvaccinated fall with the vaccination rate, up to the herd immunity threshold



Notes: Figure shows the simulated infections per unvaccinated capita, over the course of a pandemic, as a function of the vaccination rate, for various levels of infectiousness given by R_0 .

Figure E.2: A 20 percentage point increase in the vaccination rate causes a large reduction in infections among the unvaccinated, regardless of starting level



Notes: Figure shows the change in the share of the unvaccinated that are ever infected, when the vaccination rate increases by 20 percentage points, from a given initial vaccination rate for various levels of infectiousness given by R_0 .