Online appendix

Direct and indirect effects of vaccines: Evidence from COVID-19

Seth Freedman, Daniel W. Sacks, Kosali Simon, and Coady Wing

Contents

A Sensitivity Analysis

A.1 Household Indirect Effects

We explore the sensitivity of our results to four specification choices: how we adjust for pre-period covariates, the range of birth dates used to define treatment and control children, Poisson instead of OLS, and the handling of measurement error.

Our first robustness test adjusts for vaccination status and COVID incidence as of April, 2021. In Appendix Table A.6, Panel A, we reproduce our main estimates but model the propensity score as a function of vaccination and prior COVID as well as age. This has essentially no effect on our estimates. 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 range of birth dates to define treatment and control children. In Panels B and C of Appendix Table A.6, we show results for two alternative approaches, using 180 or 540 day ranges. The shorter cutoff results in a larger point estimate for the indirect effect, but also a larger standard error; the estimates are statistically indistinguishable. In Appendix Table A.4 we show estimates from Poisson regression models.

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 a bit larger than but 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.[1] 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.

A.2 School-wide Indirect Effects

Our finding of no spillovers from vaccinated schoolmates is robust to alternative samples and specifications, as we show in Appendix Table A.7 and Appendix Table A.4. In particular, we expand the sample to include students with imputed school assignment, limit the sample to include students with unique

[1] The 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.

schools (rather than just unique school type), control for instructional modality and mask mandate presence, adjust for prior infection via the propensity score (as in the direct effect design), and estimate poisson models. Across these specifications the reduced form and instrumental variables estimates are positive and insignificant, we continue to estimate school-wide spillovers that are much smaller than the family-level spillover, albeit often not statistically significantly so.

Our results are also robust to design. Our main design holds fixed students' *cohort* in the pre-period and post period, so their grade changes. Our results are not sensitive to this design choice. In our alternative design, we hold fixed the grade, so in the pre-period sample consists of 2020 sixth graders. As in our main design, the treatment group is students who will go to middle school in sixth grade, and the control group is students who will go to elementary school in sixth grade. This design better controls for school effects, but at the expense of possible grade effects (in the sense that in our main design but not this design, the treatment group changes schools from pre-period to post period). We report the results in Appendix Table A.8. The estimates are quite similar and we continue to estimate a small (but now negative, i.e. incorrectly signed) indirect effect. The IV point estimate is now positive and significantly different from the family-level estimate.[2]

[2] The first stages are identical in this design and our main design. To understand why, recall that this is a DID estimate, i.e. $(\bar{y}post, treat - \bar{y}_{post,control}) - (\bar{y}pre, treat - \bar{y}_{pre,control})$. The first difference is the 2021 difference in vaccination rates (own or school wide) among middle school versus elementary school sixth graders, the same in both designs. The second difference in our main design is the 2020 difference between the same group, where as in our alternative design it is the 2020 difference between fifth graders. However for the first stage, the outcome is vaccination, which is always zero in 2020, so the different pre-period comparison does not affect the first stage.

Table A.1: Sample restrictions and sample counts

Restriction	# People
A. Direct effects design	
Born May 13, 2008-May 12, 2009, or Nov 3, 2009-Nov 2, 2010	133,048
... and alive January 1, 2020	133,013
B. Household-level indirect effects design	
Proper age, alive Jan-1-2020, valid address, treated child or control child present	254,577
... and just treated or just control child, and no interim children	238,684
... and 8 or fewer household members	159,254
C. School-level indirect effects design	
All 6th graders, 2021-22	65,969
... and alive Jan-1-2020	65,950
... and has address	46,511
... and has school	44,562
... and school unambiguously middle or elementary	29,060
... and school not imputed	27,114
D. Grade-level indirect effects design	
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, "proper age" is 2-10 or 30 or older, and 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	-0.00010	-0.00014	0.00042
	(0.00031)	(0.00007)	(0.00013)	(0.00029)
IV estimate	0.00188	-0.00047	-0.00064	0.00199
	(0.00145)	(0.00031)	(0.00061)	(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 **??**.

Table A.3: Direct effects: Sensitivity analysis

Dep. var.	1+ dose	2+ doses	Any COVID	Non-COVID ER
A. Adjust for prior COVID via the propensity score				
DID estimate	0.2259	0.2122	-0.0021	0.0004
	(0.0015)	(0.0015)	(0.0002)	(0.0003)
IV estimate			-0.0101	0.0021
			(0.0011)	(0.0014)
Vaccine effectiveness			0.760	-0.172
			(0.024)	(0.134)
# People				133,013
B. Limit to born within 180 days of the cutoff				
DID estimate	0.2207	0.2090	-0.0028	0.0002
	(0.0022)	(0.0021)	(0.0003)	(0.0004)
IV estimate			-0.0134	0.0010
			(0.0016)	(0.0020)
Vaccine effectiveness			0.818	-0.072
			(0.024)	(0.161)
# People				64,420
C. Expand to born within 540 months of cutoff				
DID estimate	0.2290	0.2149	-0.0023	0.0003
	(0.0013)	(0.0012)	(0.0002)	(0.0002)
IV estimate			-0.0108	0.0012
			(0.0009)	(0.0009)
Vaccine effectiveness			0.757	-0.152
			(0.018)	(0.130)
# People				195,730
D. Has Address				
DID estimate	0.2524	0.2391	-0.0031	0.0005
	(0.0023)	(0.0022)	(0.0004)	(0.0004)
IV estimate			-0.0129	0.0022
			(0.0015)	(0.0018)
Vaccine effectiveness			0.793	-0.180
			(0.025)	(0.174)
# People				60,914

Notes: See notes to Table **??**. Table reports identical estimates, except the specification or sample differs as indicated. Specifically, in panel A we estimate a propensity score as a function of indicator for any COVID infection as of April, 2021, and weight by the inverse propensity score. In panel B the sample is limited to children born within 180 days of the vaccine eligibility cutoff; in C it is extended to 540 days. In Panel D it is limited to students with an address (as used in the household-level indirect effects design).

Table A.4: Poisson model estimates

Specification	DID (Reduced form)
A. Direct effect	
DID estimate	-0.350
Standard error	(0.033)
Percent effect	<-0.295>
B. Indirect effect- household	
DID estimate	-0.075
Standard error	(0.026)
Percent effect	<-0.072>
C. Indirect effect - school	
DID estimate	-0.069
Standard error	(0.065)
Percent effect	<-0.067>

Notes: Table reports coefficients estimated from Poisson models for the indicated design. The DID estimate the coefficient on $treat \times post$. The percent effect is obtained as $\exp(\beta) - 1$. See notes to Tables **??**, **??**, and **??**] for details on sample and design. Robust standard errors, clustered on individual/household/school district, in parentheses.

Table A.5: Sources of household vaccination rate increases

Group	All ages (1)	Ages 0-10 (2)	Ages 11-13 (3)	Ages 14-18 (4)	Ages 18-29 (5)	Ages 30+ (6)
DID Estimate (Child eligibility)	0.09675	-0.00009	0.08996	0.00379	0.00228	0.00081
	(0.00203)	(0.00005)	(0.00078)	(0.00068)	(0.00049)	(0.00138)
Outcome mean (Dec, 2021)	0.3195	0.0015	0.0604	0.0387	0.0222	0.1967
N Addresses	64,950	64,950	64,950	64,950	64,950	64,950
N People	159,254	159,254	159,254	159,254	159,254	159,254

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 **??**. 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 respondent (a 2-10 or 30+ year-old).

Table A.6: Household-level indirect effects: Sensitivity analysis

Outcome	Household Vaccination (1)	Own Vaccination (2)	Any COVID (3)	Severe COVID (4)
A. Adjust for prior COVID and prior vaccination				
DID Estimate (Child eligibility)	0.09735	0.00263	-0.00086	0.00001
	(0.00203)	(0.00246)	(0.00030)	(0.00012)
IV Estimate (Family Vaccination)			-0.00885	0.00008
			(0.00309)	(0.00121)
Percent effect			<-0.62650>	<0.02874>
N Addresses	64,950	64,950	64,950	64,950
N People	159,254	159,254	159,254	159,254
B. Use 180 day window for child birthday				
DID Estimate (Child eligibility)	0.09044	0.00338	-0.00099	0.00005
	(0.00279)	(0.00337)	(0.00042)	(0.00016)
IV Estimate (Family Vaccination)			-0.01095	0.00057
			(0.00457)	(0.00180)
Percent effect			<-0.77848>	<0.21331>
N Addresses	33,991	33,991	33,991	33,991
N People	83,746	83,746	83,746	83,746
C. Use 540 day window for child birthday				
DID Estimate (Child eligibility)	0.10358	-0.00055	-0.00081	-0.00001
	(0.00175)	(0.00211)	(0.00026)	(0.00010)
IV Estimate (Family Vaccination)			-0.00785	-0.00007
			(0.00249)	(0.00097)
Percent effect			<-0.56373>	<-0.02735>
N Addresses	88,654	88,654	88,654	88,654
N People	216,742	216,742	216,742	216,742
D. Obviously related instrumental variables				
DID Estimate (Child eligibility)	0.14195	0.00473	-0.00155	0.00010
	(0.00252)	(0.00310)	(0.00043)	(0.00016)
IV Estimate (Family Vaccination)			-0.01091	0.00069
			(0.00301)	(0.00110)
Percent effect			<-0.720>	<0.239>
N Addresses	56,705	56,705	56,705	56,705
N People	101,613	101,613	101,613	101,613

Notes: The sample and specification are identical to those in Table **??**, except each panel differs in one dimension. In panel A we adjust for prior COVID and prior vaccination (both as of April, 2021) via the propensity score, which we specify as a fully saturated function of indicators for age, prior COVID, and prior vaccination. 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 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.7: School-level indirect effects: Sensitivity analysis

Outcome	School Vaccination (1)	Own Vaccination (2)	Any COVID (3)
A. Include students with imputed school assignment			
DID Estimate (Middle school effect)	0.22428	0.00263	0.00104
	(0.01556)	(0.00640)	(0.00127)
IV Estimate (School vaccination rate)			0.00466
Percent effect			<0.24622>
Test $\beta_1 = -0.010$			0.012
N Students	59,345	59,345	59,345
N Districts	192	192	192
B. Limit to students with unique school assignment			
DID Estimate (Middle school effect)	0.22292	0.00498	-0.00022
	(0.01818)	(0.00756)	(0.00152)
IV Estimate (School vaccination rate)			-0.00100
Percent effect			<-0.05242>
Test $\beta_1 = -0.010$			0.207
N Students	23,912	23,912	23,912
N Districts	174	174	174
C. Control for instruction modality and mask mandate			
DID Estimate (Middle school effect)	0.21405	-0.00052	-0.00032
	(0.01627)	(0.00667)	(0.00138)
IV Estimate (School vaccination rate)			-0.00149
Percent effect			<-0.07833>
Test $\beta_1 = -0.010$			0.210
N Students	27,086	27,086	27,086
N Districts	181	181	181
D. Inverse propensity-score weighted for prior infection			
DID Estimate (Middle school effect)	0.22525	0.00252	0.00078
	(0.01620)	(0.00676)	(0.00149)
IV Estimate (School vaccination rate)			0.00347
			(0.00662)
Percent effect			<0.18340>
Test $\beta_1 = -0.010$			0.049
N Students	27,113	27,113	27,113
N Districts	182	182	182

Notes: See notes to Table **??**. 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. Because these time varying regressors are not balanced between treat and control, we also include school-district and date (year-month) fixed effects here. In panel D we adjust for prior COVID incidence via the propensity score. Each panel reports p-value of the test that the effect of school-wide vaccination rates, β_1, equals the point estimate for the effect of household vaccination rates. Robust standard errors, clustered on school district, in parentheses.

Table A.8: School-level indirect effects: Fixed grade design

Outcome	School Vaccination (1)	Own Vaccination (2)	Any COVID (3)
DID Estimate (Middle school effect)	0.22526	0.00292	0.00116
	(0.01620)	(0.00668)	(0.00134)
IV Estimate (School vaccination rate)			0.00516
Percent effect			<0.27085>
Test $\beta_1 = -0.010$			0.014
N Students	55,963	55,963	55,963
N Districts	185	185	185

Notes: See notes to Table **??**. The specification is identical but the sample consists of a fixed grade rather than a fixed cohort. Specifically the pre-period includes sixth graders in 2020-21, rather than fifth graders. Robust standard errors, clustered on school district, in parentheses.

Table A.9: 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	0.0007
	(0.0123)	(0.0013)
Bandwidth	81.2	95.9
# Observations	31,079	35,998
B. Born within 1 year of August 1, 2010		
RD Estimate	0.0014	-0.0021
	(0.0031)	(0.0014)
Bandwidth	63.9	79.6
# Observations	23,151	28,978

Notes: Table reports placebo tests for the effect of more vaccine-eligible grade mates 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: Population counts: INPC vs. US Census

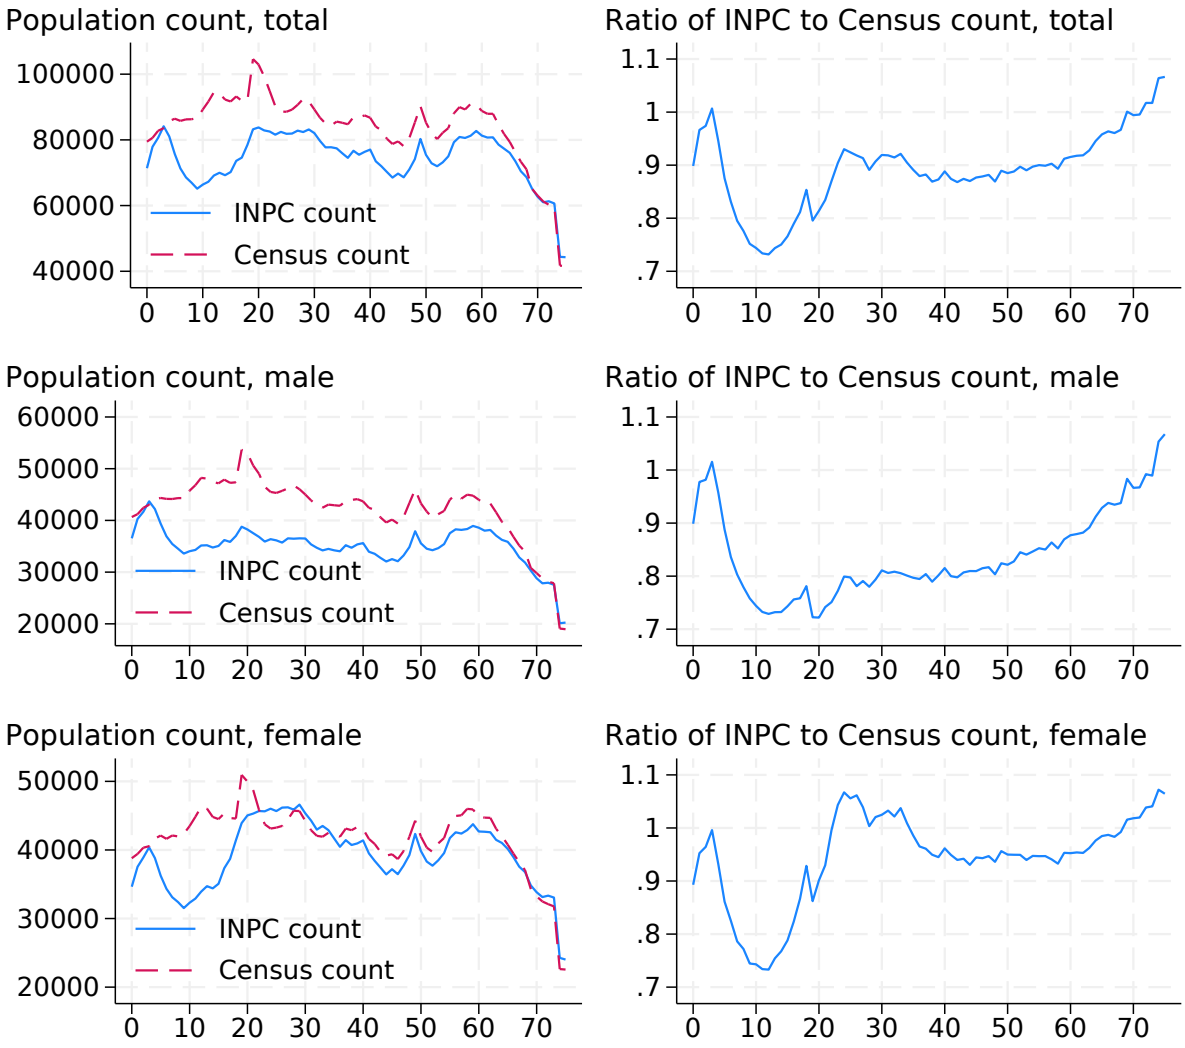

Notes: Figure reports population counts (left side) from the INPC and from Census (U.S. Census Bureaul (2024b), as of April 1, 2020), by age, overall and by sex, as well as the ratio of INPC to Census counts (right side).

Figure A.2: Comparison of Indiana and United States

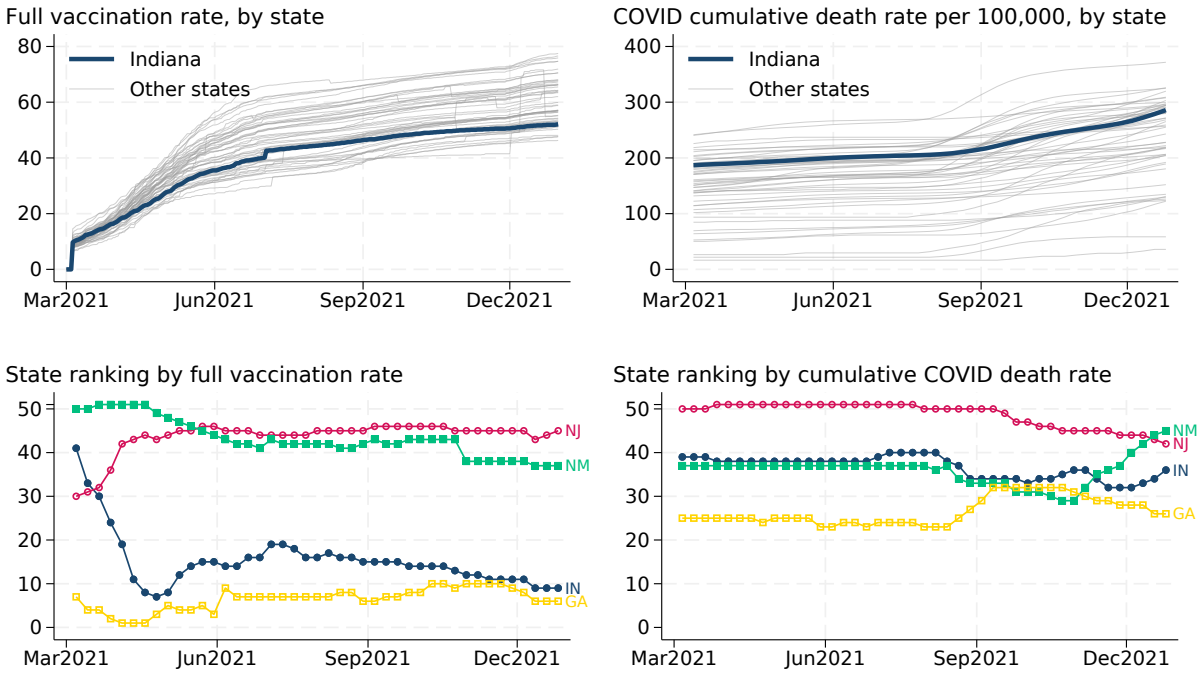

Notes: Figure comapres Indiana to other states in terms of full vaccination rates (Center for Disease Control (2024a), left side) and cumulative COVID death rate (Center for Disease Control (2024b), right side). The top panels show the time series for Indiana and each other state, and the bottom panels show the time series of Indiana's ranking among states (as well as select other states). Population counts from U.S. Census Bureaul (2024a).

Figure A.3: COVID-19 vaccine eligibility by date of birth

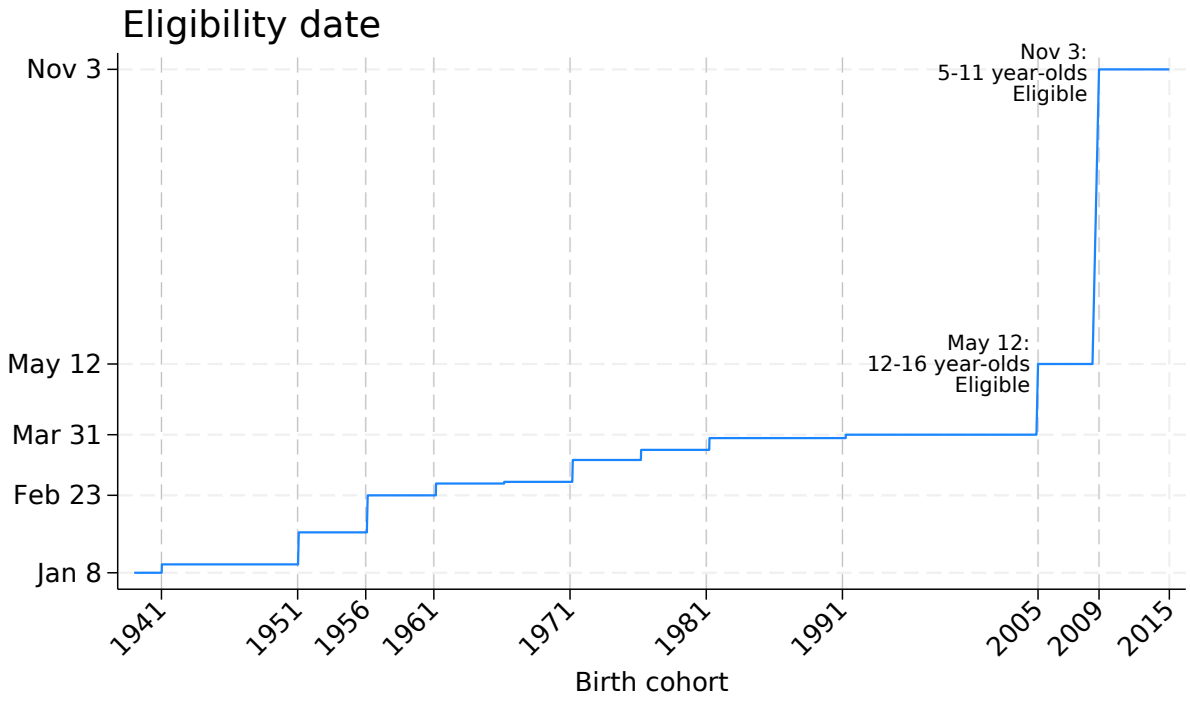

Notes: Figure shows the date (in 2021) Indiana expanded eligibility for the COVID-19 vaccines to different age groups, by day of birth.

Figure A.4: Direct effects: Inverse propensity-score weighted event study

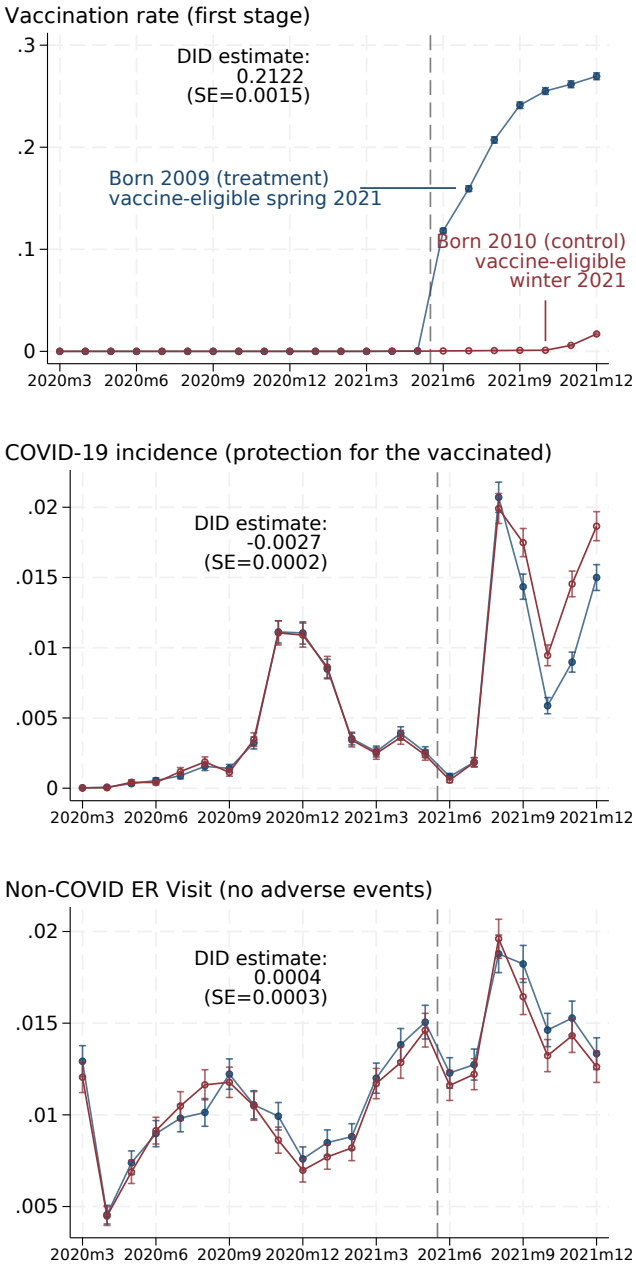

Notes: See notes to Figure **??**. This figure is identical except it is propensity score weighted. The figure shows that the level of COVID incidence matches almost exactly in the pre-period. As the propensity score is a function of April, 2021 COVID incidence, the average pre-period COVID rates will match, but the trend is not guaranteed to match. The figure therefore provides evidence to support the parallel trends assumption, conditional on the propensity score.

Figure A.5: Complier Weights

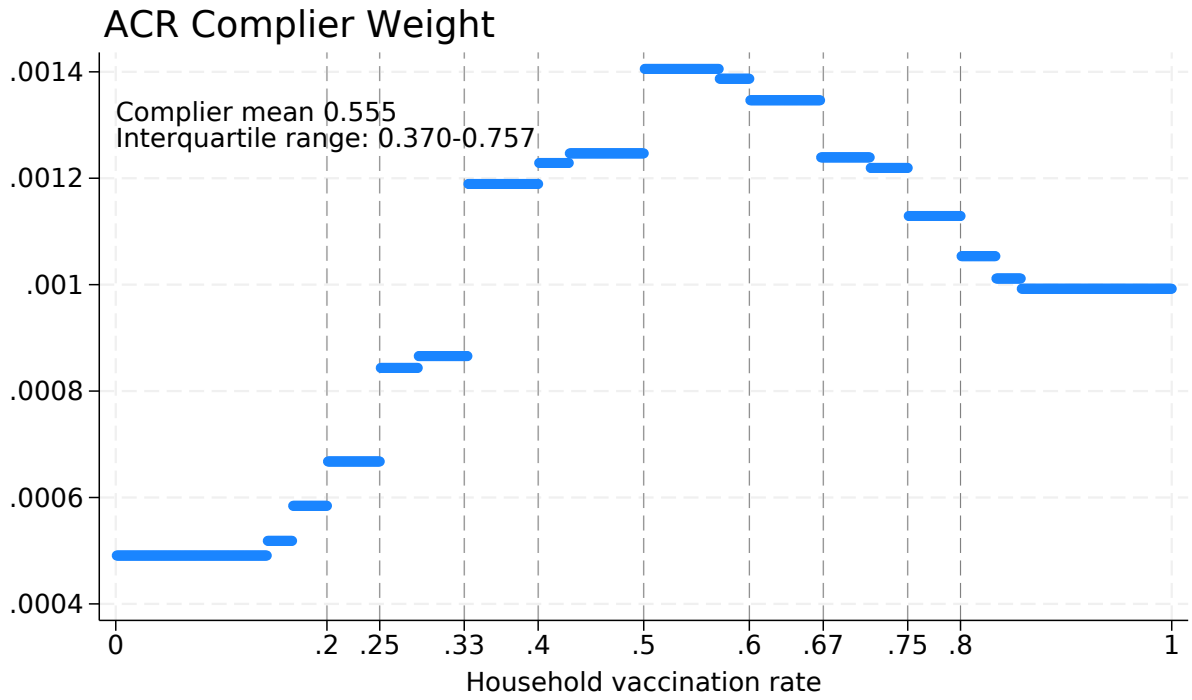

Notes: Figure shows the distribution of weights underlying the interpretation of our 2SLS estimates as a weighted-average of complier marginal effects. The weight shows the weight put by the estimator on the fraction of compliers with a control ($z = 0$) vaccination rate below the indicated level, and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

Figure A.6: Household-level indirect effects: Inverse propensity-score weighted event study

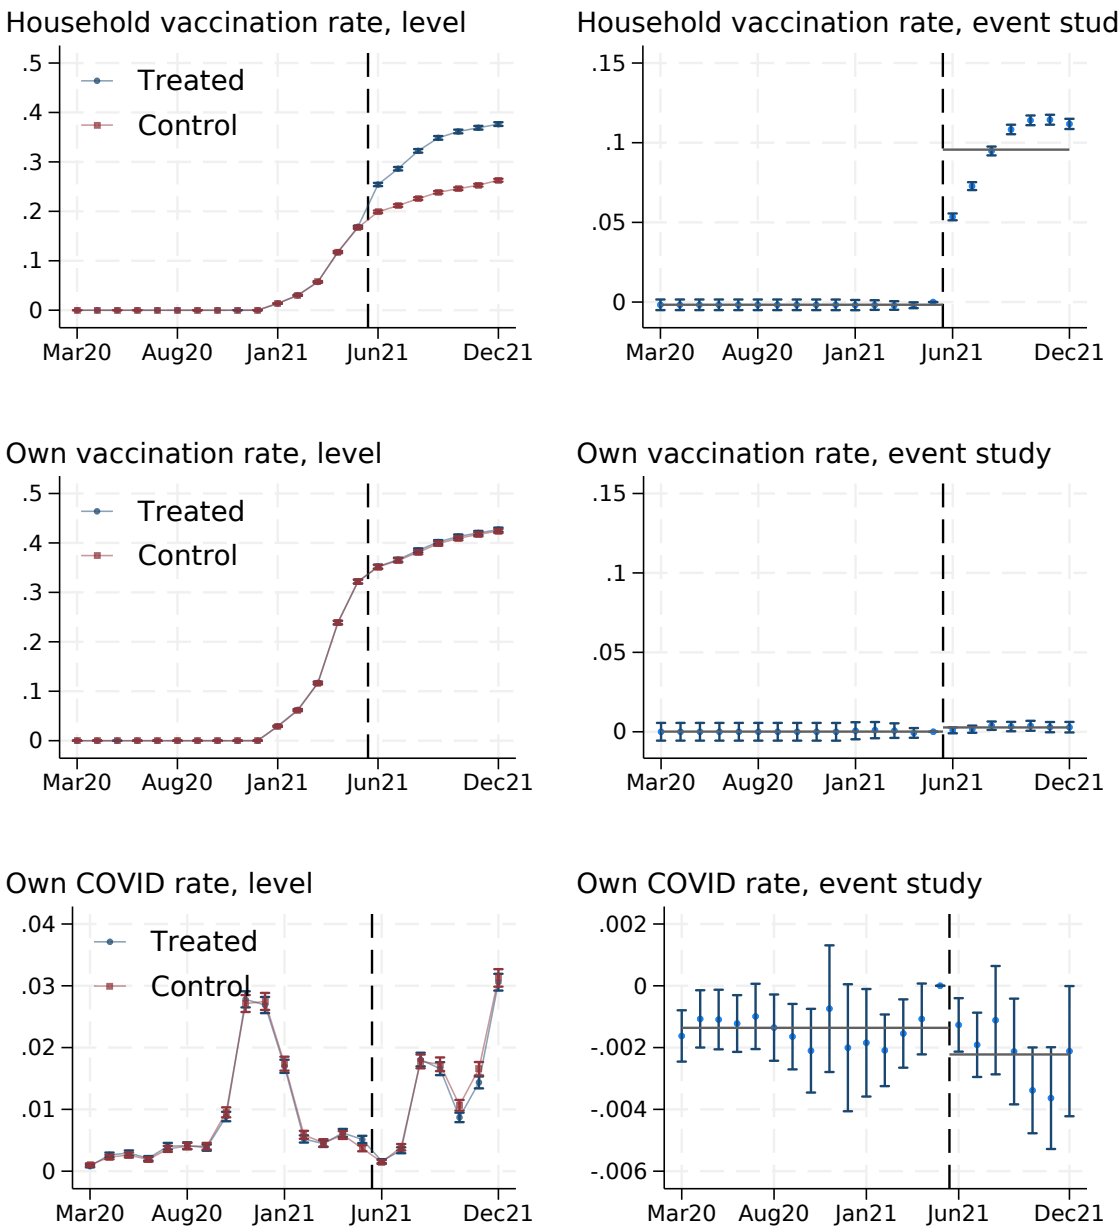

Notes: Figure is identical to Figure **??** except the propensity score includes vaccination status (as of April, 2021), and ever having COVID (as of April 2021), all fully interacted with age indicators.

Figure A.7: COVID incidence among fifth and seventh graders

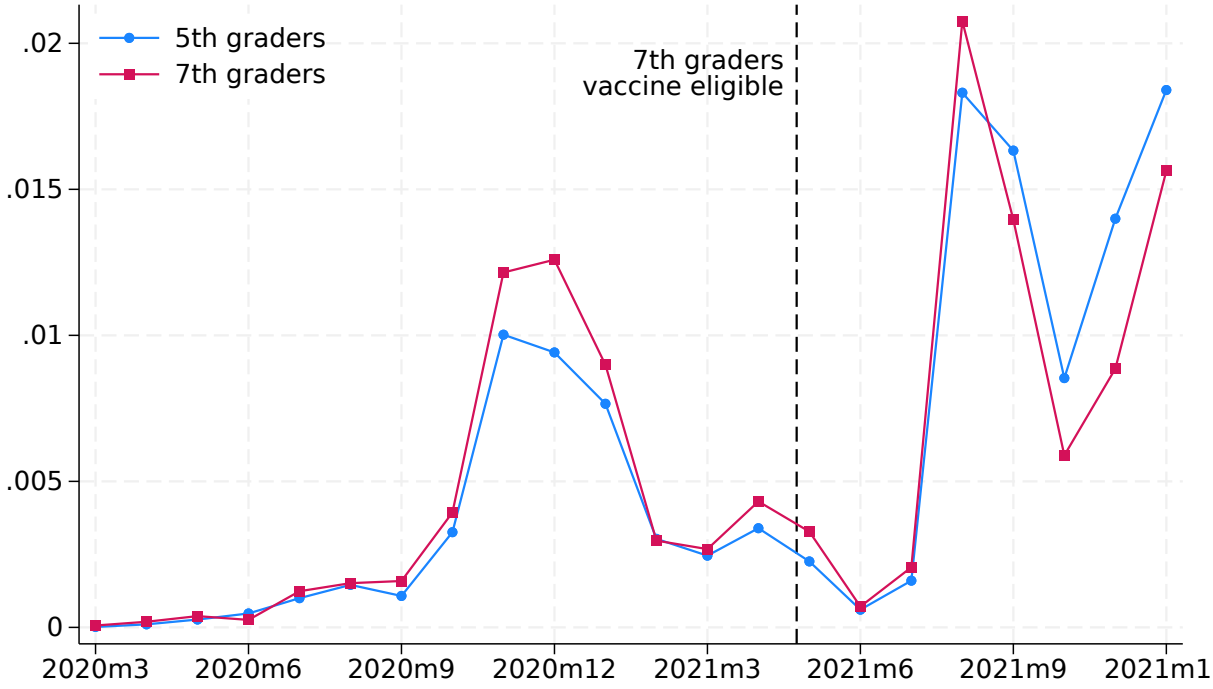

Notes: Figure shows the fraction of 5th and 7th graders with at least one positive COVID test in each month.

Figure A.8: School-level indirect effects: Inverse propensity-score weighted event study

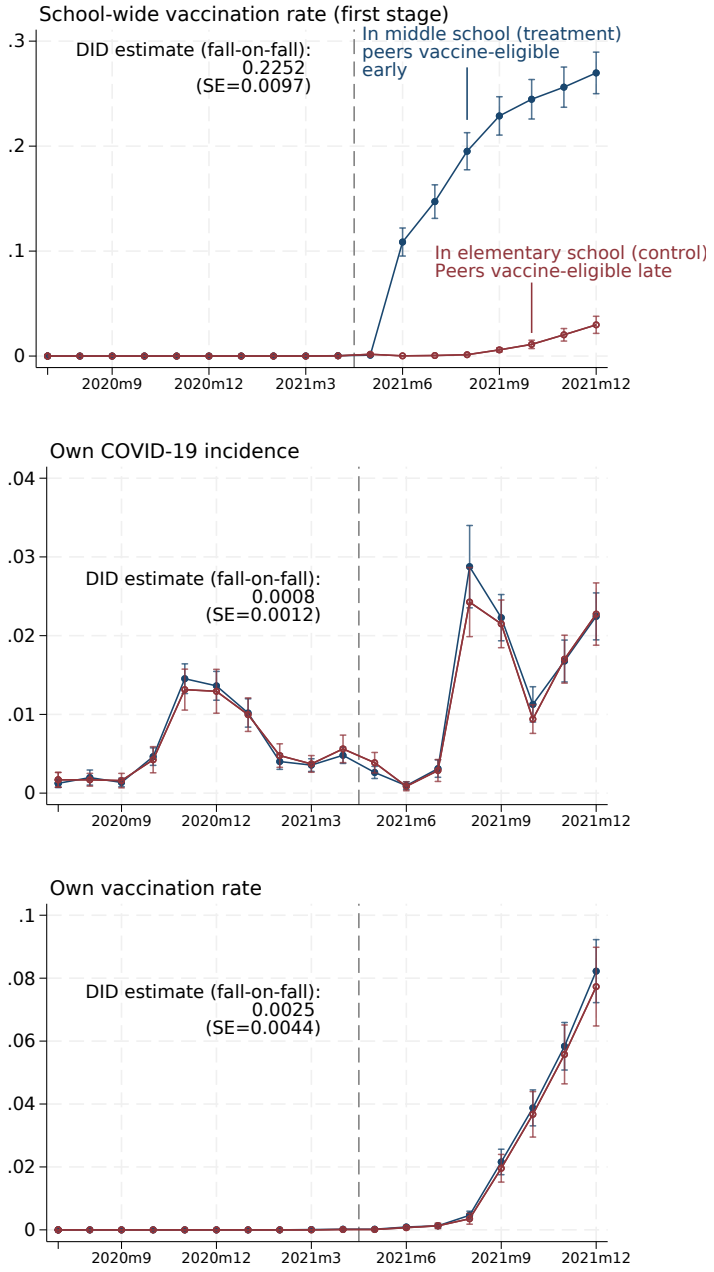

Notes: See notes to Figure **??**. This figure is identical except it is propensity score weighted. The figure shows that the level of COVID incidence matches almost exactly in the pre-period. As the propensity score is a function of April, 2021 COVID incidence, the average pre-period COVID rates will match, but the trend is not guaranteed to match. The figure therefore provides evidence to support the parallel trends assumption, conditional on the propensity score.

Figure A.9: Monthly regression discontinuity estimates

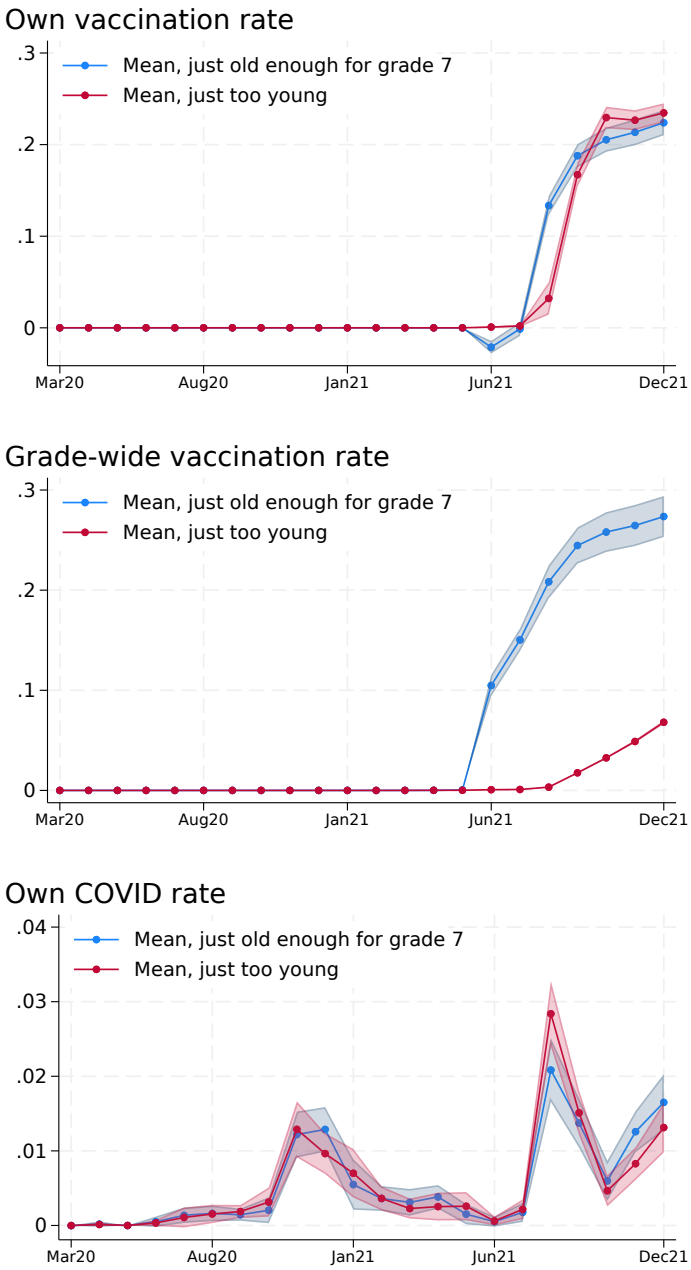

Notes: This figure plots monthly averages of the indicated outcome for children born just before and just after the threshold for school starting age (August 1, 2009). See notes to Figure **??** for details on the sample. We estimate these averages as α_0 and β_{ITT} from Equation **??**. 95% confidence based on robust standard errors, clustered on date of birth (the running variable) are in parentheses.

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 $i = 1...N$ 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 $j = 1...N$ 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 $i = 1...N$.

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 $i = 1...N$.

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 defined 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 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

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

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 **??**.

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 quality imposes the independence assumption to drop conditioning on the instrument. The fifth line decomposes the expectation 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 o 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 your 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

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

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 $Pr(test = 1)$.

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, $Pr(positive|test = 1)$, soemtimes 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, $Pr(positive|test = 1)$ 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	-0.037	-0.004	0.019
	(0.005)	(0.005)	(0.013)
Coef. on post	0.117	0.075	0.100
	(0.003)	(0.004)	(0.011)
Coef. on treat	0.013	0.001	0.016
	(0.003)	(0.003)	(0.008)
Constant	0.109	0.145	0.102
	(0.002)	(0.002)	(0.007)
# Observations	91,056	119,207	18,393
# Clusters	54,815	44,565	410
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 **??**, **??**, and **??** 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 **??** we estimate household indirect effects using two stage least squares regressions. The estimating equations we use are:

$$
\begin{aligned}
Y_{it} &= \beta_0 + \beta_1 \widehat{Vacc_{-it}} + \beta_2 treat_i + \beta_3 post_t + \epsilon_{it} \\
Vacc_{-it} &= \alpha_0 + \alpha_1 EarlyEligChildPresent_i \cdot post_t + \alpha_2 EarlyEligChildPresent_i + \alpha_3 post_t + \varepsilon_{it}
\end{aligned}
$$

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 $j = 0...1000$ different values. The values on the grid correspond to houshold 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 $v_0 = 0$, $v_1 = .001$, ... $v_{1000} = 1$.

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 is she were exposed to instrument setting $Z_i = z$ and household vaccination rate $V_{-it} = v_j$.

Impose the standard instrumental variable assumptions:

1. 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 formally, 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 $Vacc_{-it} = v_j$. 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:

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

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

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

Figure A.5 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 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.[3]

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 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.

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

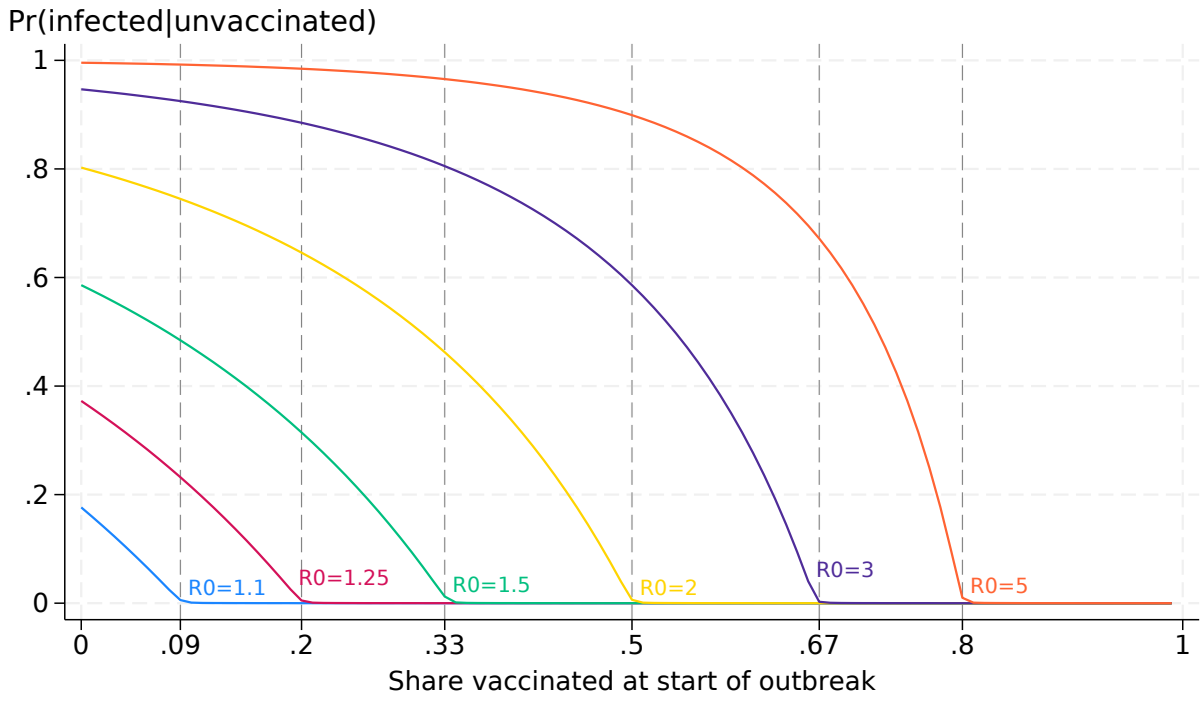

Notes: Figure shows the simulated infections per unvaccinated capita, over the course of a pandemic, as a function of the vaccination rate, for various levels of infectiousness given by R_0.

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

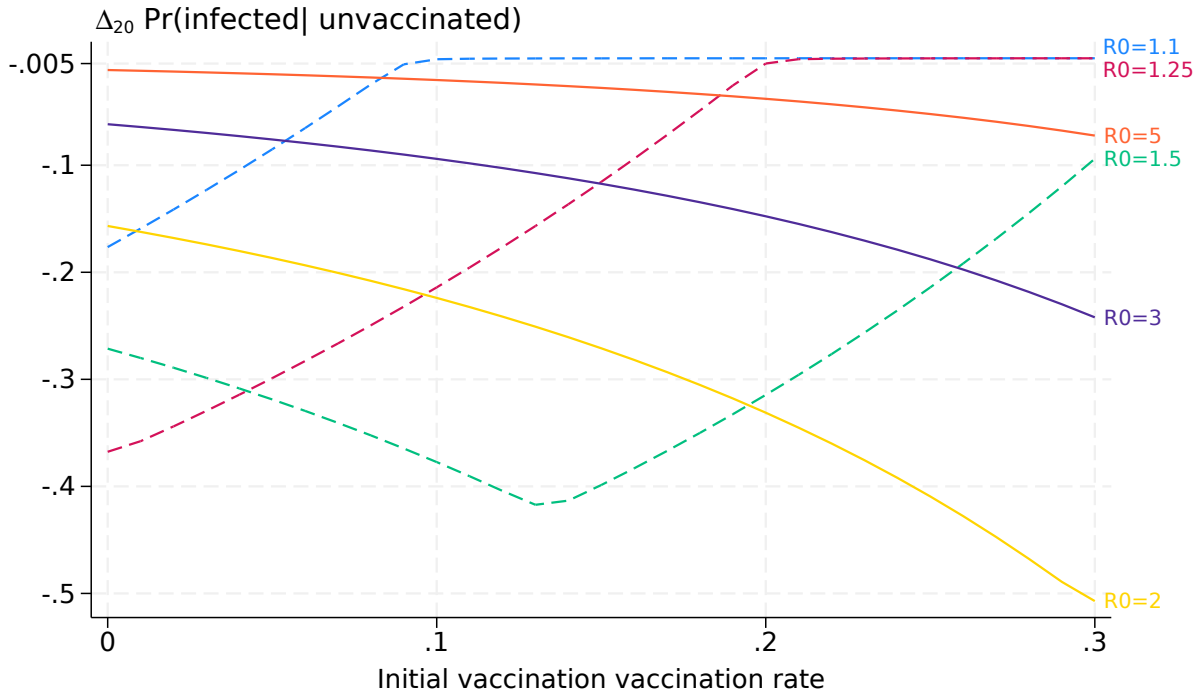

Notes: Figure shows the change in the share of the unvaccinated that are ever infected, when the vaccination rate increases by 20 percentage points, from a given initial vaccination rate for various levels of infectiousness given by R_0.

E.3 Ex Post Power Analysis

Our analysis of school-wide indirect effects is premised on the idea that schoolmates are the relevant peer group for our study population of vaccine ineligible sixth graders. One concern is that schoolmates are only *part* of the relevant peer group for our sixth graders. The true peer group is unknown but it is could be some mixture of schoolmates, grademates, and general community members. In this case, the middle school vs elementary school variation that underlies the first stage of our school-wide indirect effects design might be overestimating the shift in "true" peer vaccination rates. In addition, the first stage shift may occur at different points along the dose-response relationship between vaccination rates and outcomes, which could be non-linear.

To explore how these concerns may affect the statistical power of our study, we used the estimated standard error from our school wide indirect effects IV models to assess the (ex post) statistical power to reject the null of no indirect effect when the true indirect effect was assumed to be the value generated from the SIR model.

We examined four scenarios:

1. In scenario 1, we use our main results in which the peer group is the school wide vaccination rate.

This implies a first stage shift from 2% to 22%.

2. In scenario 2, the peer vaccination rate is comprised of half community members with a 50% vaccination rate, and half school wide vaccination rate of 22% in the treatment group and 2% in the control group). This implies a first stage shift from 26% to 36%.

3. In scenario 3, the peer vaccination rate is comprised of half grade mates with a 6% vaccination rate, and half school wide vaccination rates of 22% in the treatment group and 2% in the control group. This implies a first stage shift from 4% to 14%.

4. In scenario 4, the peer vaccination rate is an even average of community members with a 50% vaccination rate, grade mates with a 6% vaccination rate, and school wide vaccination rates with 22% in the treatment group and 2% in the control group. This implies a first stage shift from 19% to 22%.

The results of this exercise are in table E.1. In scenario 1, which corresponds to our main analysis, power is high for all but the highest value of R0. In scenario 2, the model implied "true effect" is very small under a low R0 and the power is low in these scenarios. For R0 above 1.25, the implied indirect effect ranges from .019 to .19 and power is between 90% and 100%. In scenario 3, which leads to a shift from very low vaccination to 14% vaccination, the implied treatment effects are larger and well powered for R0 below 5. In scenario 4, treatment effects are small for very low R0 and very high R0 and power is low in these cases. But power is high for the implied treatment effects for R0 between 1.25 and 3.

Table E.1: Power Analysis Results for Different Mixing Scenarios

r_0	True Effect	p-value	Power
Scenario 1: The peer group is the school.			
1.10	-0.1414	0.0000	1.0000
1.25	-0.3439	0.0000	1.0000
1.50	-0.2891	0.0000	1.0000
2.00	-0.1679	0.0000	1.0000
3.00	-0.0670	0.0000	1.0000
5.00	-0.0125	0.0265	0.6142
Scenario 2: The peer group is a mix of the community and school.			
1.10	-0.0000	0.4986	0.0504
1.25	-0.0001	0.4936	0.0517
1.50	-0.1915	0.0000	1.0000
2.00	-0.1680	0.0000	1.0000
3.00	-0.0750	0.0000	1.0000
5.00	-0.0192	0.0015	0.9084
Scenario 3: The peer group is a mix of the grade and school.			
1.10	-0.1042	0.0000	1.0000
1.25	-0.1781	0.0000	1.0000
1.50	-0.1295	0.0000	1.0000
2.00	-0.0744	0.0000	1.0000
3.00	-0.0289	0.0000	0.9977
5.00	-0.0050	0.2208	0.1907
Scenario 4: The peer group is a mix of the community, grade, and school.			
1.10	-0.0000	0.4974	0.0507
1.25	-0.0254	0.0000	0.9890
1.50	-0.1409	0.0000	1.0000
2.00	-0.0839	0.0000	1.0000
3.00	-0.0356	0.0000	0.9999
5.00	-0.0079	0.1113	0.3354

Notes: In scenario 1, treatment effectively changes the peer vaccination rate from 2% to 22%. In scenario 2, from 26% to 36%; in scenario 3, from 4% to 14%, in scenario 4, from 19% to 26%.

References

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. Rouphael, C. B. Creech, et al. (2020). Efficacy and safety of the mrna-1273 sars-cov-2 vaccine. *New England journal of medicine*.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica 82*(6), 2295–2326.

Center for Disease Control (2024a). Covid-19 vaccinations in the united states, jurisdiction [dataset]. `https://data.cdc.gov/Vaccinations/COVID-19-Vaccinations-in-the-United-States-Jurisdi/unsk-b7fc/about_data`. Last accessed May 16, 2024.

Center for Disease Control (2024b). Provisional covid-19 death counts by week ending date and state [dataset]. `https://data.cdc.gov/NCHS/Provisional-COVID-19-Death-Counts-by-Week-Ending-D/r8kw-7aab/about_data`. Last accessed May 16, 2024.

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.

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.

Imbens, G. W. and J. D. Angrist (1994). Identification and estimation of local average treatment effects. *Econometrica 62*(2), 467–475.

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.

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*.

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.

U.S. Census Bureaul (2024a). Natioanl population totals: 2020-2023 [dataset]. https://www.census.gov/data/tables/time-series/demo/popest/2020s-national-total.html. Last accessed June 11, 2024.

U.S. Census Bureaul (2024b). State population by characteristics: 2020-2023 [dataset]. https://www.census.gov/data/tables/time-series/demo/popest/2020s-state-detail.html. Last accessed June 11, 2024.