National Institute Economic Review (2021), 257: 1, 14–35
doi:10.1017/nie.2021.23

NIER CAMBRIDGE UNIVERSITY PRESS

ARTICLE

THE GREAT COVID-19 VACCINE ROLLOUT: BEHAVIOURAL AND POLICY RESPONSES

M. Christopher Auld[1]* and Flavio Toxvaerd[2,3]

[1]Department of Economics, University of Victoria, Victoria, BC, Canada
[2]Faculty of Economics, University of Cambridge, Cambridge, United Kingdom
[3]Centre for Economic Policy Research, London, United Kingdom
*Corresponding author. Email: auld@uvic.ca

(Received 21 April 2021; revised 22 June 2021; accepted 1 July 2021)

Using daily data on vaccinations, disease spread and measures of social interaction from Google Mobility reports aggregated at the country level for 112 countries, we present estimates of behavioural responses to the global rollout of COVID-19 vaccines. We first estimate correlates of the timing and intensity of the vaccination rollout, finding that countries which vaccinated more of their population earlier strongly tended to be richer, whereas measures of the state of pandemic or its death toll up to the time of the initial vaccine rollout had little predictive ability after controlling for income. Estimates of models of social distancing and disease spread suggest that countries which vaccinated more quickly also experienced decreases in some measures of social distancing, yet also lower incidence of disease, and in these countries, policy-makers relaxed social distancing measures relative to countries which rolled out vaccinations more slowly.

JEL codes: I12; C5.

Keywords: economic epidemiology; econometrics; COVID-19; vaccination.

1. Introduction

Widespread vaccination against the virus that causes COVID-19 disease began in December 2020, and as of the end of March 2021, almost 600 million doses had been distributed.[1] Concerns have been raised that as vaccinations progress, people will reduce their social distancing, blunting the efficacy of vaccination in reducing the spread of the disease (e.g. Galanti *et al.,* 2021). To the best of our knowledge, this is the first paper attempting to estimate the effects of vaccinating a population on social distancing.[2]

Vaccinations against COVID-19 disease have been shown to substantially reduce the probability that a vaccinated person will exhibit symptomatic infection.[3] Holding behaviour constant, epidemiological models of disease spread predict that more vaccinations must, then, reduce new infections. However, economic models predict that both vaccinated and unvaccinated people will change their risky behaviours as the proportion of the population which is vaccinated increases (Chen and Toxvaerd, 2014). Depending on their pattern, these behavioural responses may act to reinforce the protective effects of vaccination, or may reduce the efficacy of vaccination in reducing new infections. Determining their net effect is an empirical problem.

[1]Authors' calculation from *Our World in Data* dataset described in Section 3.

[2]Several studies use simulations and data from one country to estimate the effect of vaccinations on disease dynamics, taking behaviour as invariant to own-vaccination status and population vaccination rates: see for example Cook and Roberts (2021); De-Leon *et al.* (2021); Public Health England (2021a).

[3]Evidence on effectiveness compiled as of March 2021 is presented in Public Health England (2021b). See also Forni and Mantovani (2021) for a review of the development and biological efficacy of COVID-19 vaccines.

We estimate the effects of an increase in vaccinations per capita on various behavioural measures drawn from Google Mobility reports at the country level, exploiting data on the outcomes of 112 countries across the world. These data are limited in the sense that we only observe country-level averages, so that, among other aggregation issues, we cannot distinguish between changes in average behaviour resulting from changes in the behaviours of vaccinated and unvaccinated people. However, global data are inherently interesting in the sense that we are able to quantify outcomes for the majority of the world's population, and the staggered rollout of vaccinations across countries allows us to estimate the effects of vaccination by comparing countries which initiated vaccinations sooner and more intensely to those that vaccinated later or less intensely.

The most challenging econometric issue arises because vaccinations were not randomly assigned across countries. To the extent that countries which vaccinated sooner did so because they were experiencing more severe COVID-19 epidemics—or, even more problematically, were expecting more severe epidemics in the future when vaccine contracts were written in the Fall of 2020—statistical models may mistake vaccinations as a cause of behaviour with vaccinations as a consequence of behaviour. That is, we may observe that there are more vaccinations in countries with more severe epidemics because the world successfully triaged vaccinations, not because vaccinations caused disease. Lacking randomisation in vaccinations, we cannot fully resolve this issue. We specify panel data models which mitigate these concerns with various fixed effects and other panel data strategies, consider a variety of specifications of these models to assess robustness, and present estimates of the determinants of the rollout of vaccinations across countries in an effort to assess how the pattern of vaccination rollouts across countries was related to the severity of their epidemics.

We find that countries which vaccinated more of their populations earlier did tend, unconditionally, to be countries experiencing more severe epidemics in late 2020. But, this association essentially disappears once income is held constant: each doubling of GDP per capita was associated with approximately six more vaccinations per capita by the end of our sampling window (19 March 2021), whereas measures of the severity of the epidemic as of or up until the end of November 2020 no longer have any substantial nor statistical ability to predict the vaccination rollout. At least initially, vaccines tended to be rolled out faster in richer, not sicker, countries. It is disturbing that countries in which the most lives could have been saved by vaccination did not, it appears, tend to receive vaccinations first, but these results do lend credibility to estimates of the effects of vaccination on behaviour: factors other than the state of the pandemic primarily drove the initial vaccination patterns.

Our major results include: some measures of social distancing, notably retail visits, fell in quickly-vaccinating countries relative to statistically comparable countries which were not able to vaccinate as quickly, but at the same time, new cases fell in relatively quickly vaccinating countries. We also find evidence that restrictions intended to increase social distancing were relaxed, relatively, in countries in which vaccines were rolled out quickly, suggesting that policy makers view vaccinations and restrictions as substitutes. We conclude by discussing important limitations to these results.

2. Analytical framework

Conventional epidemiological models do not feature endogenous behavioural responses to the risk of infection. Hence, a standard model of the effect of vaccinating some fraction of the population against COVID-19 either would not consider behavioural responses, or would impose behavioural response in an *ad hoc* fashion.[4] In contrast, economic epidemiological models focus on how endogenous changes in behaviour alter spread of disease and the impact of interventions designed to limit spread.[5]

[4]Examples include Ferguson *et al.* (2020) and Iboi *et al.* (2020).

[5]See Philipson (2000) for a survey of the early literature in economic epidemiology which arose largely as an attempt to understand aspects of the HIV/AIDS epidemic. Fenichel *et al.* (2011) provides an overview of the literature on multiple diseases up to that time. McAdams (2020) discusses some of the contributions from economic epidemiology to the study of COVID-19. Toxvaerd and Rowthorn (2020) contrast the economic approach to vaccination to the traditional public health approach,

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

A long-standing concern in the economic epidemiology literature revolves around the possibility that vaccinations will cause offsetting behavioural responses which reduce the efficacy of the vaccine in reducing infections or, if behavioural responses are large enough, even perversely cause more infections.[6] This literature focusses on the case in which an imperfect prophylactic vaccine reduces a vaccinated person's incentive to reduce costly efforts to avoid infection. People who are not vaccinated, however, may also change their behaviour in response to an increase in the proportion of people around them who are vaccinated.

To fix ideas, consider a very simple, static rational choice model in which a susceptible agent chooses a contact rate c to maximise

$$U(c) = B(c) - \gamma Rc - p\left(c, I, V, \overline{V}\right)L\left(\overline{V}\right), \tag{1}$$

where $U(c)$ denotes the utility function, $B(c)$ denotes the benefits the agent obtains from a contact rate of c, R denotes a measure of the stringency of governmental restrictions on social interactions, $p(\cdot)$ denotes the perceived probability of infection which depends on the contact rate, the proportion of the population that is infected, I, the agent's own vaccination status denoted by an indicator V (which we treat as exogenous) and the proportion vaccinated $\overline{V}$, and L denotes the loss, in units of utility, of becoming infected. To make the model as simple as possible, assume the functional forms for $B(\cdot)$, $L(\cdot)$ and $p(\cdot)$ that permit us to write,

$$U(c) = log\,(c) - \left[\gamma R - \left(\beta I - \alpha V - \delta\overline{V}\right)\left(\pi_0 - \pi_1\overline{V}\right)\right]c, \tag{2}$$

where all parameters are positive. Here, β is a transmission coefficient which depends on biology and social structure, α represents how the individual perceives that her own vaccination status affects her probability of infection, δ how she perceives population vaccination levels to affect her own probability of becoming infected, and π_0 represents the loss due to infection for a narrowly self-interested agent. The loss due to infection decreases at rate π_1 as population vaccinations increase, reflecting altruism: particularly since countries tended to prioritise the most vulnerable in distributing vaccines, the agent may perceive that the cost to becoming infected is lower as vaccinations are rolled out because they are less likely to infect and harm or kill other people through their own behaviour.

Assuming the agent ignores any effect her own outcomes have on the population vaccination or infection rates, the model predicts that she will choose contact rate

$$c^*\left(I, V, \overline{V}\right) = \frac{1}{\gamma R + \left(\beta I - \alpha V - \delta\overline{V}\right)\left(\pi_0 - \pi_1\overline{V}\right)}. \tag{3}$$

Assume that parameters, vaccination rates and infection rates are such that each term in parentheses in the denominator on the right-hand side is positive. Then it is easy to show that the model predicts that, other things being equal, more stringent restrictions, higher disease prevalence, a higher transmission probability, the agent herself not being vaccinated or a higher overall loss to infection (π_0) reduce risky behaviour. The effect of the agent's vaccination status on her probability of becoming infected, though, is ambiguous: if she selects a sufficiently higher contact rate in response to her own vaccination, the total effect will be an increase in her chances of infection.[7] Vaccinating more of the population around this individual also unambiguously increases her risky behaviour; this effect occurs both because she perceives her own risk to be lower at the margin, and because the costs of becoming infected decrease.

whereas Makris and Toxvaerd (2020) consider how the anticipated arrival of vaccines alters incentives to engage in social distancing.

[6]See for example Auld (2003); Geoffard and Philipson (1997); Talamas and Vohra (2018); Toxvaerd (2019).

[7]That is, if response is sufficiently elastic, then $p\left(c^*\left(I, V=1, \overline{V}\right), I, V=1, \overline{V}\right) > p\left(c^*\left(I, V=0, \overline{V}\right), I, V=0, \overline{V}\right)$.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

Integrating such behavioural responses into models of disease propagation shows that they will generally spur the epidemic relative to models in which behaviour is fixed, although not necessarily in models with heterogeneity or at very high levels of risk (Kremer, 1996).

With these behavioural considerations in mind, we consider an econometric implementation of this simple model in Section 4.

3. Data

3.1. Google Mobility reports

Our main measures of social distancing behaviour are drawn from Google Mobility data. These data and more detail on their construction can be found at: https://www.google.com/covid19/mobility/. The version of these data we use for analyses was downloaded 12 April 2021. Using Google Maps searches, these data track visits and lengths of stay to locations categorised as Retail and Recreation, Grocery and Pharmacy, Parks, Transit stations, Workplaces and Residential places.

We consider all five of these categories as measures of social mobility: *Retail and Recreation* provides a measure of visits to restaurants, cafes, shopping centers, theme parks, museums, libraries and movie theaters. *Workplaces* data provide information on visits to places of work. *Transit stations* measures use of public transport hubs such as subway, bus and train stations. *Parks* tracks visits to 'places like local parks, national parks, public beaches, marinas, dog parks, plazas and public gardens'. Finally, *Residential* data provide information on visits to places of residence; we sometimes refer to this outcome as 'time at home'. We interpret increases in all of these except residential visits as decreases in social distancing, and increases in residential visits as an increase in social distancing.

According to Google, the location accuracy in these data varies by country, but the country fixed effects we include in all models should mitigate this concern. Google reports these data as changes relative to the 5-week period from 3 January to 6 February 2020. We reconstruct the absolute variables up to a constant by calculating one plus the log of the value reported by Google.

3.2. Vaccinations and measures of epidemic progress

To the best of our knowledge, the only comprehensive international data on vaccinations have been compiled by the *Our World in Data* team.[8] We use the smoothed data on new vaccinations; the version of these data we use in the analyses was downloaded on 11 April 2021. These data were constructed from a wide variety of official sources, most commonly the Ministry of Health or similar government entity. They do not include vaccinations given during clinical trials. Note that we select number of vaccinations given as our the covariate of interest in all analyses rather than the proportion vaccinated or fully vaccinated: we measure the effect of an additional vaccination, regardless of whether the arm into which it is injected has already received at least one vaccination. We choose this outcome because the data on the number of people vaccinated contain more missing observations.

We also use daily new COVID-19 cases and deaths per capita attributed to COVID-19 as compiled from official sources and reported and smoothed by *Our World in Data.*

3.3. Other data

The Oxford COVID-19 Government Response Tracker (OxCGRT) policy stringency index is used as a composite measure of policy responses intended to reduce COVID-19 infections. This index can in

[8]The data, a variety of descriptive statistics, and more information on the construction of the data can be found at: https://ourworldindata.org/covid-vaccinations.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

principle vary between zero and one, and is constructed as a simple average of dummy variables indicating a variety of restrictions imposed with the intention of increasing social distancing, such as school closures and restrictions on gatherings. See Hale *et al.* (2021) for further information on the construction of this index.

We obtain information on temperature and humidity from the Air Quality Open Data Platform Worldwide COVID-19 dataset, and take the arithmetic average of readings from multiple stations within a country on a given day.[9] The version of these data we use in the analysis was downloaded 11 April 2021. These data do not provide comprehensive global coverage. As we demonstrate below, including weather in our models has essentially no effect on the estimates, so we do not include weather in most of our analysis so as to retain as large a sample as possible.

Data on GDP per capita were drawn from the *Our World in Data* dataset.

3.4. Sample selection

We face a tradeoff in determining the date on which to begin the estimation sample. The vaccination rollout begins in December 2020. Using data from before the rollout provides more information to estimate the parameters other than vaccination effects, and provides a 'pre-treatment' period which in effect generates more control groups, but as we shift the first date in our analyses backwards in time we also implicitly impose more assumptions on the temporal stability of the parameters we estimate, and we place more weight on pre-treatment outcomes as controls. We choose 1 December 2020 to begin the estimation sample; the estimates are not sensitive to small changes in this date. The last date we use in the estimation was 19 March 2021, as, due to reporting delays, missing data starts to become a substantial concern after this date. We, then, have 109 days in our estimation sample.

We do not include countries which have no data on vaccinations. These may be countries for which vaccination has occurred but the data are not yet reported, or they may be countries which have not begun vaccinations. In either case, they provide a dubious control group for vaccinating countries. We also exclude countries with no Google Mobility data. After these exclusions, we are left with an estimation sample of 112 countries in the main analyses. We list these countries along with means of selected outcomes in Appendix Table A3. Notably, we cannot use China or India in our analyses due to missing data, so a substantial fraction of the world's population is missing.

Three countries (Angola, Honduras and Uganda) have a total of four missing observations on the Google Mobility data during the sampling window, and there are a total of 19 observations with zero reported daily new cases. Since we model the logarithm of new cases, we drop these observations. The number of observations included in the models then varies slightly (between 12,185 and 12,204) depending on whether lagged variables are included and on whether the dependent variable is new cases or one of the Google Mobility outcomes.

Appendix Table A1 presents descriptive statistics for the key outcomes in our analyses.

4. Econometric methods

4.1. Main models

Consider estimating the simple model of behaviour above, equation (3), which says that a person's behaviour is a function of their own vaccination status, population vaccinations and disease prevalence. Aggregating across people, the model implies that changes in vaccination rates across populations will affect average behaviour for several reasons: people may change their behaviour when they are vaccinated, they may respond to population vaccination rates given their own status for altruistic reasons, and vaccination affects epidemic dynamics which in turn affect behaviour. We take this model

[9]These data and more information about them are available from: https://aqicn.org/data-platform/covid19/.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

to the data by deploying standard panel econometric models for various behavioural outcomes. Our least restrictive specification takes the form,

$$y_{it} = \beta_0 + \beta_1 V_{i,t-14} + \beta_2 y_{i,t-14} + \mathbf{X}_{it}\gamma + \alpha_i + \varphi_t + \lambda_i t + u_{it}, \quad (4)$$

where i indexes countries and t indexes time in days: y_{it} denotes a behavioural outcome or new cases of disease, V_{it} denotes vaccinations per capita, $\mathbf{X}_{it}$ is a vector of controls, including policy stringency and in some models weather conditions, α_i are country effects, φ_t are day effects, $\lambda_i t$ are country-specific linear trends and u_{it} is a disturbance term.

The effects of vaccinations on various behaviours, the parameters β_1, are identified in this model by variation in the timing of the rollout of vaccinations across countries. The country fixed effects remove all variation in both behaviours and vaccination patterns that can be attributed to time-invariant or predetermined outcomes, such as the severity of the epidemic up to the beginning of the vaccine rollout and socioeconomic development. The inclusion of a complete set of day effects sweeps out common time trends without imposing any functional form assumptions, and subsume seasonal effects, day of the week effects and holiday effects. Controlling for weather, as measured by temperature and humidity, removes the concern that we may mistake changes in behaviour and vaccinations that both result from, for example, a particularly inclement day with an effect of vaccinations on behaviour.[10] Including a lagged value of the dependent variable flexibly controls for nonlinear trends in the dependent variable and, in models in which disease incidence is the outcome, helps reduce bias resulting from countries increasing vaccinations in response to worsening COVID-19 conditions.

The effects of policies designed to reduce social interaction on disease spread are the subject of considerable research.[11] Controlling for policy stringency implies that we estimate the effects of vaccinations which do not result from changes in policies which themselves are caused by, or correlated with, the vaccination rollout. We assess possible problems with estimating and intepreting the causal effect of vaccinations if policy stringency is also affected by vaccinations below. As the focus of this paper is not to measure the effects of policy on behaviour, we do not attempt to estimate the distinct effects of each of the policy responses from which this index was constructed, but since these estimates may themselves be of interest, we report them along with the estimated effects of vaccinations.

Finally, the country-specific linear trends relax the conditionally parallel trends assumption on which identification would otherwise rest. Much like previous research in other contexts (e.g. Lee and Solon, 2011; Rambachan and Roth, 2019), we find that the estimates are sensitive to this assumption; standard two-way fixed models in some cases lead to different qualitative conclusions than models with country-specific trends. Since identification in the models without trends is based on the more restrictive and more implausible assumption that behavioural outcomes across the world would have trended identically if not for the vaccination rollout, we prefer the model with trends.

In models with a lag of the dependent variable, we ignore the endogeneity of this variable due to our relative large-T sample of 109 days. In some models, we also lag the covariate of most interest, vaccinations, and the policy stringency index, by 14 days. We do so to modestly mitigate concerns of simultaneity bias, allow time for vaccinations to be effective, and to allow time for both policies and vaccinations to affect behavioural outcomes, particularly reported new cases.[12]

We estimate standard errors and other measures of sampling variability using the method of Driscoll and Kraay (1998), with up to 7 days serial dependence (the estimates are not sensitive to this choice of lag). Inference is then robust to general forms of heteroskedasticity, serial correlation, and cross-sectional dependence, mitigating concerns that spatial or temporal dependence bias our standard errors.

[10] See Mecenas *et al.* (2020) for a survey of research on COVID-19 and weather.

[11] See for example Berry *et al.* (2021); Chen *et al.* (2020); Goolsbee and Syverson (2021); Wright *et al.* (2020); Yan *et al.* (2021), among many others.

[12] See Bellemare *et al.* (2017) for a discussion of the limitations of lagging the covariates as general solution to simultaneity concerns.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

In a variety of robustness checks, we assess the sensitivity of these results to various assumptions. Note, however, that our two-way fixed models with many time periods and staggered interventions are subject, even with arbitrary country trends, to the criticism levelled by Goodman-Bacon (2018), Callaway and Sant'Anna (forthcoming), and other recent papers that such models implicitly use inappropriate control groups and recover at best variance-weighted averages of heterogeneous treatment effects. However, to the best our knowledge, methods to circumvent these issues when the treatment is continuous have not yet been developed. Restricting our sample to countries which have started vaccinations by the end of our sampling period may reduce this bias.

4.2. Determinants of vaccinations and simultaneity bias

Note that we do not control for time-varying measures of epidemic severity due to the potential of simultaneity bias. The estimates of the effects of vaccines should then be considered reduced forms that reflect both the direct effects of vaccination and indirect effects occurring due to changes in disease spread. The same concern applies to vaccination rates, and our estimates of the effects of vaccination rates are biased to the extent that governments procured more vaccinations during the initial rollout if they anticipated a worsening of the epidemic at the same time. Such a mechanism implies that vaccinations are endogenous in model (4) due to simultaneity, that is, that vaccinations may 'reverse cause' behaviour and disease incidence. However, as discussed in Section 5.1, the pattern of the initial rollout was largely determined by contracts written months earlier, suggesting that this may not be a major concern.

To provide suggestive evidence on the extent to which the pattern of rollouts reflected disease severity, we present estimates of models investigating the determinants of the rollout of vaccinations across countries. We construct a cross-sectional dataset of country characteristics and two measures of epidemic severity as of early December 2020, immediately preceding the initiation of vaccinations: the total of number of COVID-19 deaths (per capita) up to that time, and the mean number of new cases per day (per capita) during November 2020. We use total COVID-19 deaths per capita up until 1 December 2020 as a measure of the cumulative effects of the epidemic on the country up to that date, and the number of new cases per day as a measure of the severity of the epidemic immediately prior to the initial rollout of vaccines in the following month. Standard errors in these cross-sectional models are robust only to heteroskedasticity.

5. Results

5.1. The rollout of vaccinations across nations

Why did some nations rollout vaccinations faster than others? The panel data models we present below hinge on the assumption that the variation remaining in the speed of the vaccination rollout, after sweeping out fixed effects and covariates, is exogenous. That is, some nations distributed more vaccinations than others for reasons unrelated to those affecting how risky behaviours would have evolved in the absence of a vaccine rollout. In this section, we present some suggestive evidence that at least the initial rollout of vaccines was largely unrelated to the severity of the epidemic in each country, which in turn suggests that the regression-adjusted correlations we present between vaccinations and the outcomes we consider largely reflect the consequences of vaccinations, not the causes of vaccinations.

The institutional background is consistent with the proposition that the world largely failed to triage vaccinations. Vaccinations began to be distributed in December 2020 based on contracts between manufacturers and governments signed months earlier, which implies that subsequent changes in the severity of a country's epidemic likely had little or no effect on the magnitude of its rollout. Media reports suggest the rollout was largely determined by income, by political savvy and idiosyncratic pharmaceutical approval processes, and by the happenstance of the locations of manufacturers. For example, the *New York Times* reported on 15 December, the week of the initial rollout, that 'rich countries have

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

"cleared the shelves" of vaccines' (Twohey *et al.*, 2020). After noting that India had secured a relatively large number of doses from Novavax, the *Times* article quotes the CEO of that company explaining why that occurred: 'India gets priority because it's my home country'.

The lag between government contracts for vaccine purchases and delivery also suggests that the global pattern of the vaccine rollout could not have been largely determined by variation in epidemic severity in late 2020 and early 2021. Even by the end of August 2020, government contracts for more than one billion doses had been signed, and by December 2020 developed countries had preordered about 10 billion doses more.[13] Governments and charitable organisations have, however, also contracted with manufacturers in an effort to provide more timely access to vaccinations, notably through COVID-19 Vaccines Global Access (COVAX). But, this initiative had not, by the end of our estimation window (19 March 2021), provided a substantial number of doses, and aims to reach only about 3.3 per cent of the total population of each participating country by the end of the first half of 2021 (Gavi: The Vaccine Alliance, 2021).

Consider first the correlates of the timing and intensity of the rollout of vaccinations across nations. In Table 1, we present estimates of a variety of models of total vaccinations by 19 March 2021.

Column (1) of Table 1 presents estimates of the simplest specification including only our measure of epidemic severity immediately prior to the beginning of the vaccine rollout, mean cases per day during November 2020. The statistically significant estimate ($t = 3.8$) suggests that, unconditionally, countries with more severe epidemic tended to distribute more vaccinations during the initial rollout, although the effect is quite small: each doubling of new cases in November is associated with about $(0.013)(0.69) \approx 0.01$ more vaccinations per capita, that is, 1 more vaccination per 100 people.

Column (2) of Table 1 adds income as a control. The unconditional association between vaccinations and epidemic severity essentially disappears after controlling for income, as the estimated coefficient falls to only 0.003 with an associated *t*-statistic of 0.7. But income strongly predicts vaccinations: the point estimate suggests each doubling of GDP per capita is associated with about $(0.05)(log\,2) \approx 0.04$, or 4 more vaccinations per 100 people. Put another way, the results suggest that if country A's epidemic was worse than country B's immediately prior to the rollout, country A did tend to vaccinate faster than B, but if country A and B have the same average income and country A's epidemic was worse than B's, they vaccinated at similar rates. It is income, not disease severity, which drove the initial rollout pattern, according to these estimates.

Column (3) adds continent effects to the model, thereby comparing countries only to their continental neighbours. Using only within-continent variation has little effect on the results. The point estimate on income increases in magnitude modestly to 0.06 and remains highly statistically significant, whereas the estimated coefficient on new cases in November remains small and statistically insignificant. We report the estimated continent effects since they may be of interest: there is little evidence that, conditional on income and epidemic severity, countries on different continents vaccinated at substantially different rates.

Finally, in column (4) we add an additional measure of COVID-19's damage: total deaths attributed to COVID-19 as of 1 December 2020. The point estimate is a little less than 0.4, suggesting a moderately strong relationship between deaths and vaccinations, but it is very statistically insignificant ($t = 0.86$, $p = 0.39$). Hence, neither deaths to date, nor the contemporaneous severity of the epidemic as measured by new cases, predict the rollout of vaccinations.

Taken together, these results suggest that the rollout of vaccinations across nations favoured the richest countries, not the countries experiencing the most severe epidemics given their incomes. However, the lack of association between the state of the epidemic and the rollout of vaccines does lend some suggestive credibility to the estimates to follow of the effects of vaccinations on behaviour, as they do not suggest that the pattern of vaccinations is consistent with vaccinations being triaged to the hardest-hit populations.

[13] Authors' calculation from contract listing provided by the UNICEF COVID-19 Vaccination Market Dashboard, available from https://www.unicef.org/supply/covid-19-vaccine-market-dashboard. See also So and Woo (2020) for an analysis of purchase commitments as of mid-November 2002.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

Table 1. Correlates of vaccinations across countries

	(1)	(2)	(3)	(4)
log (new cases during November)	0.013***	0.003	0.006	0.007
	(3.76)	(0.69)	(1.11)	(1.37)
log (GDP per capita)		0.053***	0.063**	0.047***
		(4.40)	(3.23)	(3.89)
Log (deaths by Dec 1)				0.384
				(0.86)
Asia			−0.027	0.015
			(−0.58)	(0.62)
Europe			−0.064	−0.023
			(−1.02)	(−0.49)
North America			−0.048	−0.013
			(−0.95)	(−0.34)
Oceania			−0.102	−0.050
			(−1.44)	(−1.25)
South America			−0.061	−0.039
			(−1.19)	(−0.97)
N	145	139	145	145

Note: Ordinary least square (OLS) estimates using country-level cross-section. The dependent variable is vaccinations per capita as of 19 March 2021. The omitted continent is Africa. New cases in are expressed as a mean per 100 capita per day and deaths as total deaths to date per million capita. Parentheses report *t*-statistics based on robust standard errors. Asterisks denote statistical significance at the 0.05, 0.01 and 0.001 levels.

5.2. Effects of vaccinations on behaviours and disease incidence

We vary the specification of the model across the panels of Table 2 to assess sensitivity. In all models, the outcomes are in logs, vaccinations are per capita, and the policy stringency index is bound on the unit interval. An estimated coefficient on vaccinations, then, can be interpreted as the predicted effect of increasing vaccinations from zero to a number equal to the entire population (which, since we measure vaccinations rather than people vaccinated, does not mean the entire population is vaccinated), whereas the coefficient on policy stringency can be interpreted as the predicted effect of a country 'switching on' all policies tracked by OxCGRT.

First consider Panel A, which presents two-way fixed effects estimates of day *t* vaccinations and policy stringency on day *t* outcomes. The estimates suggest policy stringency is highly contemporaneously correlated with more social distancing: that is, fewer visits to workplaces, retail, transit and parks, but more residential visits. These estimates are all substantively large and highly statistically significant. However, policy stringency is strongly contemporaneously associated with more new cases, which likely reflects causation from new cases to policy stringency rather than a perverse effect of policy. Vaccinations in this model are statistically significantly associated with more transit and parks visits, but also more residential visits, and not associated statistically nor substantially with new cases. These estimates, however, assume that effects of policy occur instantly, which is particularly implausible for new cases, which take time to incubate and detect. Further, these estimates are identified off the assumption that all countries' outcomes would trend in parallel if it were not for changes in the covariates.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

Table 2. Estimates of panel data models of behaviours and new cases

	(1)	(2)	(3)	(4)	(5)	(6)
	Workplaces	Residential	Retail	Transit	Parks	New cases
Panel A: Contemporaneous two-way fixed effects						
Vaccinations	0.064	0.069***	0.099	0.194**	0.452***	0.143
	(0.95)	(5.60)	(1.39)	(3.04)	(7.33)	(0.44)
Stringency index	−0.575***	0.170***	−0.937***	−0.802***	−0.501***	2.321***
	(−13.22)	(21.33)	(−20.28)	(−19.53)	(−12.56)	(11.20)
N	12,204	12,204	12,204	12,204	12,204	12,189
Panel B: Covariates lagged 14 days						
Vaccinations	0.119	0.045**	0.294**	0.316***	0.601***	−0.296
	(1.39)	(2.71)	(3.11)	(3.77)	(7.70)	(−0.72)
Stringency index	−0.097*	0.069***	−0.284***	−0.207***	−0.101*	−0.106
	(−2.05)	(7.57)	(−5.41)	(−4.46)	(−2.33)	(−0.47)
N	12,200	12,200	12,200	12,200	12,200	12,185
Panel C: Addition of country-specific linear trends						
Vaccinations	0.123	−0.089*	1.350***	0.571**	1.118***	−2.591***
	(0.59)	(−2.38)	(6.08)	(2.99)	(6.43)	(−3.54)
Stringency index	−0.001	0.030**	−0.039	−0.089	−0.022	−0.502**
	(−0.02)	(2.93)	(−0.64)	(−1.72)	(−0.46)	(−2.61)
N	12,200	12,200	12,200	12,200	12,200	12,185
Panel D: Addition of lagged dependent variable						
Vaccinations	0.112	−0.118**	1.091***	0.383*	0.725***	−2.035**
	(0.53)	(−3.21)	(4.88)	(2.02)	(4.07)	(−3.26)
Long-run effect	0.128	−0.183	1.566	0.548	1.001	−4 175***
$\widehat{\beta}_1/\left(1-\widehat{\beta}_2\right)$	(0.169)	(−0.194)	(2.306)**	(0.751)	(1.327)	(14.522)
Stringency index	0.051	−0.021*	0.208***	0.114*	0.084	−1.437***
	(0.89)	(−2.13)	(3.40)	(2.20)	(1.73)	(−8.57)
N	12,200	12,200	12,200	12,200	12,200	12,173

Note: All outcomes are in logs. Vaccinations are per capita and the policy stringency index is bound on the unit interval. Each panel reports a separate specification. The panel consists of 112 countries and 109 days, with some missingness. Complete sets of day and country effects are included in all models; panel titles note additional structure. Parentheses report t-statistics based on Driscoll–Kraay standard errors robust to arbitrary serial and spatial correlation. Asterisks denote statistical significance at the 0.05, 0.01 and 0.001 levels.

We begin to relax these assumptions in the estimates presented in Panel B of Table 2, which also deploys two-way fixed effects, but lags both vaccinations and policy stringency by 2 weeks. The estimates are sensitive to this change: the estimated effects of policy stringency on the Google Mobility behaviours both increase and decrease across outcomes, and particularly of note, the positive association between

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

new cases and policy stringency disappears: the point estimate flips sign and is not statistically significant. We interpret this result as confirmation that the apparently perverse effect reported in Panel A does reflect the endogeneity of policy to disease spread. The effects of vaccinations in this model are similar to those in Panel A.

We then relax the parallel trends assumption in Panel C by including country-specific secular trends, such that identification is achieved by estimating deviations around these trends as the vaccination rollout occurred. The inclusion of these trends has substantial effects on the estimates. Policy stringency statistically significantly increases residential visits, although the effect is rather small in magnitude: a 0.1 increase in policy stringency increases residential visits by 0.3 per cent ($t = 2.9$). The point estimates on the other Google Mobility outcomes also suggest stringency increases social distancing, although they are not statistically significant. The estimated effect of stringency on new cases suggests policies are effective at reducing disease spread: a 0.1 increase in the stringency index causes about 5 per cent lower incidence ($t = 2.6$).

The estimates of vaccinations in Panel C also differ markedly from the more restrictive estimates in Panels A and B. These estimates strongly suggest people responded to the rollout of vaccinations by reducing social distancing: distributing vaccinations equal to 10 per cent of a country's population leads to: a 0.9 per cent decrease in residential visits ($t = 2.4$), a 13.5 per cent increase in retail visits ($t = 6.1$), a 5.7 per cent increase in transit visits ($t = 3.0$) and a 11.2 per cent increase in visits to parks ($t = 6.4$). Finally, more vaccinations lead to diminished disease spread: distributing vaccinations equal to 10 per cent of a country's population reduces new cases by a very large 26 per cent ($t = 3.5$), according to this model.

Panel D presents the results of the least restrictive of our specifications, which includes both country-specific trends and a 2-week lag of the dependent variable. Here, interpretation of the estimates becomes more difficult as the point estimates reflect short-run effects, whereas long-run effects can be calculated by plugging the estimating coefficients into $\frac{\beta_1}{1-\beta_2}$, in the notation of equation (4), that is, the estimated effect over one minus the estimated coefficient on the lag of the dependent variable. These long-run effects, with inference based on the Delta method, are also displayed in the table, but as we do not think that the epidemic can reasonably be modelled as fluctuating around a long-run steady state equilibrium, we interpret them with skepticism. Qualitatively, the results in Panel D for both vaccinations and policy stringency are very similar to those in Panel C.

5.3. Parameter stability over time

The concern that our estimates are biased by vaccinations being routed to countries in which they are most needed may become more compelling as time passes. Vaccinations rerouted to areas with epidemic flare-ups bias our estimates: this mechanism will make it appear that vaccinations cause new cases (or are less effective than they are at reducing new cases) and, given that people respond to outbreaks by increasing social distancing, make it appear that vaccinations increase social distancing (or reduce social distancing less than they actually are). Further, the linear trends assumption which helps identify our models becomes more tenuous as the sampling window increases.

To investigate the stability of our estimates over time, we estimated rolling regressions. Always starting 1 December 2020, we estimate the model in Table 2 Panel C (with day and country effects and country-specific trends), ending the sample on 1st February, then, separately, on 2nd February, and so on, until 19th March. We note that changes in the estimates as we increase the sample size do not necessarily indicate that the causal effect of vaccinations is itself changing over time, as we would expect estimates to change with sample size due to sampling variability.[14] But if we interpret estimates only using earlier data as less biased than those also using later data, then we wish to assess whether we would draw qualitatively different conclusions depending on when we end the sampling window.

[14]See for example Zanin and Marra (2012) on the interpretation of rolling regression estimates.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

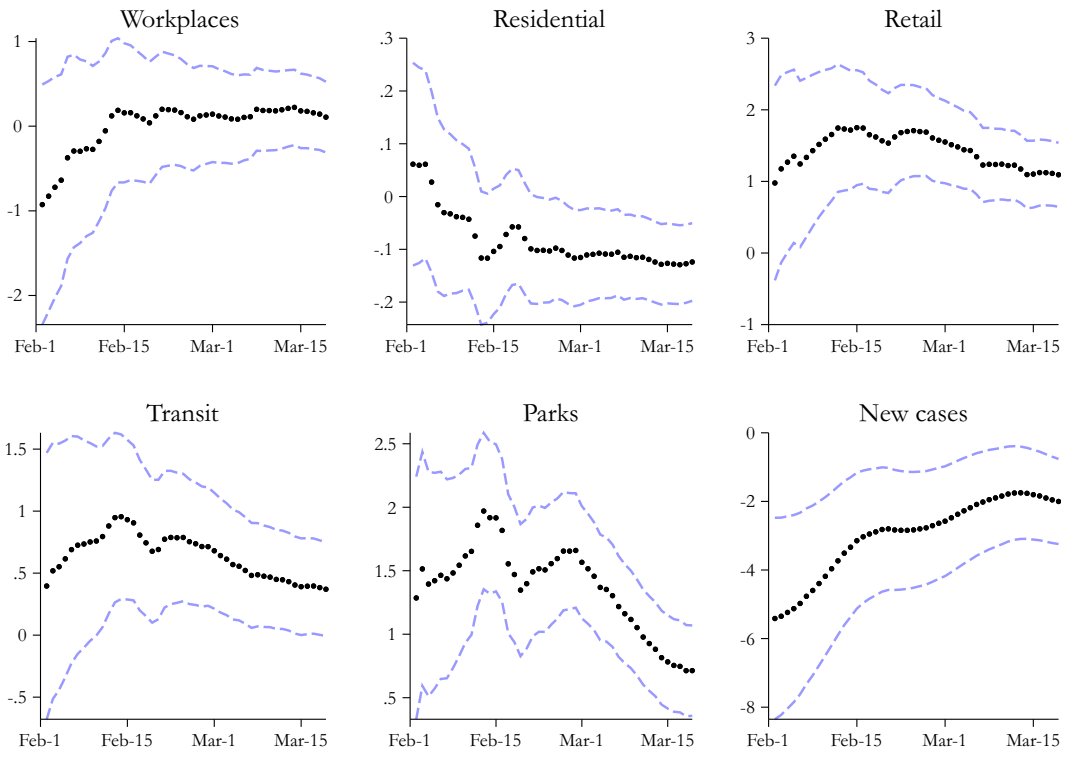

Figure 1. (Colour online) Rolling regression estimates. Figure shows each model estimated in Table 3, Panel C (models include country-specific time trends along with complete sets of day and country effects) estimated from the beginning of the sampling window (1 December 2020) to the date shown on the *x*-axis. Note the scale of the *y*-axis varies across subfigures. The dashed blue lines are 95 per cent confidence intervals

The rolling regression results are displayed in Figure 1. Many of the results change little as we increase the sampling window: the estimated effect on workplace visits diminishes in magnitude but is never statistically significant, whereas the effects on retail and transit visits change little. However, the effect of vaccinations on residential visits is initially very small and positive, but switches sign and increases in magnitude as the sampling window increases, only obtaining statistical significance when we use almost the entire sample. For all outcomes except new cases, the qualitative results are largely unaltered by varying the sampling window.

The effect of vaccinations on new cases is sensitive to when we end the sampling window. Vaccinations appear to have a very large effect on new cases during the early rollout, but this effect diminishes markedly in magnitude over time. A researcher using data from 1st December to 1st February would conclude that a 10 percentage point increase in vaccinations per capita decreases new cases by 60 per cent ($t = 4.4$). Conversely, a researcher using data from 1st December to 19th March would draw the conclusion we present in Table 2 Panel C: this effect falls in magnitude to −26 per cent ($t = 3.5$). Figure 1 demonstrates that this decline in magnitude is monotonic as the sampling window is increased. This result is consistent with increasing simultaneity bias resulting from increasingly successful vaccine triage, or with a declining causal effect of vaccinations. A declining effect may in turn result from successful within-country allocation of vaccines to people more likely to become infected, such that as an increasing proportion of the population is vaccinated, the marginal vaccinated person's reduction in infection probability decreases.

5.4. Heterogeneous effects

In Table 3, we investigate whether the effects of vaccinations differed with observed country characteristics, simply by interacting vaccinations with our measures of policy stringency and GDP per capita.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

Table 3. Heterogeneity across policy stringency and income

	(1)	(2)	(3)	(4)	(5)	(6)
	Workplaces	Residential	Retail	Transit	Parks	New cases
Vaccinations	0.250	−0.100*	1.185***	0.635**	1.015***	−1.072
	(1.00)	(−2.23)	(4.44)	(2.77)	(4.86)	(−1.22)
Stringency index	−0.000	0.030**	−0.049	−0.095	−0.031	−0.487*
	(−0.00)	(2.93)	(−0.80)	(−1.81)	(−0.63)	(−2.50)
Vaccinations × Stringency index	0.101	−0.041	0.438	0.425	0.382	0.829
	(0.19)	(−0.42)	(0.75)	(0.86)	(0.84)	(0.45)
Vaccinations × GDP per capita	−0.012	0.001	0.016	−0.006	0.010	−0.141**
	(−0.91)	(0.44)	(1.13)	(−0.47)	(0.91)	(−3.11)
N	12,200	12,200	12,200	12,200	12,200	12,185

Note: Ordinary lease square (OLS) estimates of models of Google Mobility behavioural outcomes and new COVID-19 cases. All outcomes are in logs and covariates are lagged 14 days. Complete sets of day and country effects and country-specific linear time trends are included in all models. Parentheses report *t*-statistics based on Driscoll-Kraay standard errors robust to arbitrary serial and spatial correlation. Asterisks denote statistical significance at the 0.05, 0.01 and 0.001 levels.

For ease of interpretation, the continuous covariates are first demeaned, such that the coefficients on the main effects should be interpreted as the effect at the global mean of each continuous covariate, and the coefficients on interactions should be interpreted as how this effect itself varies with the interacted regressor.

Evaluated at sample means, the estimates in Table 3 display a similar pattern as the main estimates discussed above: vaccinations decrease both social distancing and new cases, although the effect on new cases is not statistically significant. There is little evidence that vaccinations are more or less effective in changing behaviour or new cases in countries with more stringent policies, as none of the interactions between policy and vaccinations are statistically significant. Similarly, only one of the interactions between vaccination and income are statistically significant: column (6) in Table 3 suggests that vaccinations are more effective in reducing new cases in richer counties. An increase in GDP per capita of $1,000 USD decreases the estimated effect of vaccination by 0.14 units ($t = 3.1$), relative to the estimated effect at the mean of −1.06 units. That is, in wealthy countries, a given increase in vaccinations per capita decreased new infections more than in poorer countries.

5.5. Robustness

We further investigate the robustness of these results by varying the lag length on vaccinations. Instead of somewhat arbitrarily choosing 2 weeks, we vary the lag length from one to 30 days, and estimate the dynamic model with country-specific trends for each outcome at each lag length. The resulting 180 regression coefficients are presented in graphical form in Appendix Figure A1. The results of this exercise indicate that the major results discussed above are reasonably robust to lag length. The estimates of the effect of vaccination on workplace visits never obtain statistical significance, and we would draw the same qualitative conclusions over the effects of vaccinations on the remaining outcomes of interest regardless of lag length. That said, the magnitude of the effect of vaccination on residential visits increases as we consider longer lags, whereas the magnitude of the effect on new cases is largest for lags of between about 10 and 20 days.

As another robustness check, we consider whether controlling for weather conditions would affect our results. Weather may affect both disease transmission directly, may affect vaccinations, and may

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

affect social distancing, so omitting weather may lead to a spurious relationship between behaviour and vaccinations. We do not include weather as a covariate in all models because of missing data. In Appendix Table A2, we estimate the specification with country-specific trends and a lag of the dependent variable on the subset of 79 countries for which we have weather (as measured by temperature and humidity) data. We restrict attention to residential visits and new cases as key behavioural outcomes. In columns (1) and (3), we estimate models including weather controls, and in columns (2) and (4), we present versions of the models presented in Table 2 but restricted to the sample for which we have weather data. The results clearly show that the weather outcomes statistically significantly predict both behaviour and new cases, but controlling for weather has essentially no effect on the estimated coefficients on vaccinations, which change only in the third decimal place.

5.6. Effect of vaccinations on policy stringency

Do policy-makers relax restrictions as more vaccinations occur? We estimate equation (4) with policy stringency as the outcome of interest to investigate this question. As with the models of behavioural outcomes, we include complete sets of country and day effects, and we include country-specific time trends. We consider two specifications, one including a 2-week lag of the dependent variable, and one without this lag.

Table 4 presents the results. The estimate in column (1) suggests that there is a large and highly statistically significant effect of vaccinations on policy stringency. The point estimate means that distributing a number of vaccines equal to 10 per cent of a country's population decreases policy stringency by about 0.06 units ($t = 7.7$). Put another way, the descriptive statistics displayed in Appendix Table A2 show that the average stringency observed globally over our sampling window was about 0.62 and ranged between 0.12 and 0.94. The estimate in column (1) of Table 4 suggests that, if the most severely restrictive country distributed a number of vaccinations equal to its population, then its policy makers could be expected to decrease stringency to 0.34 (0.94–0.6), decreasing its stringency from the most severe to one of the least severe. Column (2) of Table 4 adds lagged stringency as a control, but the results are very similar.

These effects are both substantively and statistically significant and comprise suggestive evidence that the rollout of vaccinations would have had larger effects on disease transmission if policy stringency had not (relatively) declined in quickly vaccinating countries. That policy stringency responds to the timing of the vaccination rollout also presents a potential problem for our main estimates presented in Table 4: if policy stringency is on the causal path from vaccinations to behavioural outcomes and disease incidence, 'collider bias' may afflict our estimates of the effects of vaccination (see e.g. Pearl, 2013). We assessed the

Table 4. Effect of vaccinations on policy stringency

	(1)	(2)
	Stringency index	Stringency index
Vaccinations	−0.619***	−0.530***
	(−7.73)	(−6.75)
Stringency index (lagged 14 days)		0.219***
		(10.54)
N	12,200	12,200

Note: The dependent variable is the policy stringency index, on the unit scale. Complete sets of day and country effects are included in all models, along with country-specific linear trends. Parentheses report *t*-statistics based on Driscoll–Kraay standard errors robust to arbitrary serial and spatial correlation. Asterisks denote statistical significance at the 0.05, 0.01 and 0.001 levels.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

sensitivity of the estimates presented in Table 2 to this mechanism by re-estimating each model without including policy stringency as a covariate. The estimates on the effects of vaccination change little, suggesting that this form of bias is not a serious concern.

6. Conclusions

We conclude by summarising the main results and discussing their limitations. The evidence suggests that the rollout of vaccinations across nations resulted in:

1. A decrease in average social distancing, particularly so for visits to retail establishments.
2. Decreased new COVID-19 cases, perhaps particularly so during the earliest stages of the rollout and in relatively high-income countries.
3. Relaxation of government social distancing restrictions.

However, these results are subject to a number of important limitations.

First, while we present suggestive evidence that the rollout cannot be predicted by COVID-19 vaccinations in the month preceding the beginning of the rollout once income is held constant, and we deploy a battery of fixed effects and other panel data approaches to minimise concerns over endogeneity and uncontrolled common causes of disease spread and vaccination response, our estimates are nonetheless biased to the extent that the vaccination rollout responded to rapidly changing epidemic conditions across countries. In particular, to the extent that vaccinations are a response to changes in disease rather than a cause, our estimates will be biased in the direction of finding that population vaccination is less effective than it actually is in reducing disease spread. We find that the apparent effect of vaccinations falls as we increase the end of the sampling window from the beginning of February to late March, which, among other possible explanations, is consistent with increasing bias from increasingly successful vaccine triage across nations.

Second, we are only able to observe country-level averages, which presents a number of challenges. We cannot distinguish between the behavioural effects of vaccination on the vaccinated, and effects of population vaccination on a given person conditional on that person's vaccination status. So, for example, we have no way of disentangling a decrease in social distancing in response to more rapid vaccine distribution resulting from altruistic or narrowly self-interested motives. Moreover, since we observe only average behaviour, we cannot estimate any effects on the distribution of behaviour within countries, such as an increase in the dispersion of social distancing.

Third, we found that our estimates are sensitive to the specification of our models. Models with only contemporaneous outcomes lead to quite different conclusions than those relating current behaviour or disease incidence to policies and vaccination prior to 2 weeks. Models with only country-specific linear trends sometimes lead to different conclusions than models without these trends. We focus attention on models preferred for theoretical reasons, that is, models with lags of the covariates and which relax the parallel trends assumption, but highlight that these choices matter.

Finally, in this paper, we have sought to increase understanding of the interaction between vaccinations, policy and behaviour. We have not studied the reasons why governments have rolled out vaccines at the pace that they have. This pace will in practice be constrained by vaccine availability and by a country's ability to deliver vaccines, which in turn depends on factors such as health care systems, demographics and geography (McKee and Rajan, McKee and Rajan, 2021). Data on vaccine stockpiles are difficult to come by, as governments are reluctant to share this information publicly. In turn, vaccine delivery capacity is difficult to measure. We therefore leave these issues for further research.

We conclude that overall there is some evidence in the country-level data that countries which successfully vaccinated more quickly experienced both less social distancing and diminished spread of disease as a result of these vaccinations. We suggest future research using microdata, when they become available, which may be able to shed more light on these questions.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

Acknowledgements. We thank participants at the 'What can "economics" and "epidemiology" learn from each other?' workshop held by the National Institute of Economic and Social Research 12 February 2021, for many helpful comments.

References

Auld, M.C. (2003), 'Choices, beliefs, and infectious disease dynamics', *Journal of Health Economics*, **22**, pp. 361–77.

Bellemare, M.F., Masaki, T. and Pepinsky, T.B. (2017), 'Lagged explanatory variables and the estimation of causal effect', *The Journal of Politics*, **79**, pp. 949–63.

Berry, C.R., Fowler, A., Glazer, T., Handel-Meyer, S. and MacMillen, A. (2021), 'Evaluating the effects of shelter-in-place policies during the COVID-19 pandemic', *Proceedings of the National Academy of Sciences*, 118, p. e2019706118.

Callaway, B. and Sant'Anna, P.H. (forthcoming), 'Difference-in-differences with multiple time periods', *Journal of Econometrics*.

Chen, F. and Toxvaerd, F. (2014), 'The economics of vaccination', *Journal of Theoretical Biology*, **363**, pp. 105–17.

Chen, M.K., Zhuo, Y., Da La Fuente, M., Rohla, R. and Long, E.F. (2020), 'Causal estimation of stay-at-home orders on sars-cov-2 transmission,' preprint, arXiv:2005.05469.

Cook, T. and Roberts, J. (2021), 'Impact of vaccination by priority group on UK deaths, hospital admissions and intensive care admissions from COVID-19', *Anaesthesia*, **76**, pp. 608–16.

De-Leon, H., Calderon-Margalit, R., Pederiva, F., Ashkenazy, Y. and Gazit, D. (2021), 'First indication of the effect of COVID-19 vaccinations on the course of the COVID-19 outbreak in Israel', medRxiv.

Driscoll, J.C. and Kraay, A.C. (1998), 'Consistent covariance matrix estimation with spatially dependent panel data', *Review of Economics and Statistics*, **80**, pp. 549–60.

Fenichel, E.P., Castillo-Chavez, C., Ceddia, M.G., Chowell, G., Parra, P.A.G., Hickling, G.J., Holloway, G., Horan, R., Morin, B., Perrings, C., Springborn, M., Velazquez, L. and Villalobos, C. (2011), 'Adaptive human behavior in epidemiological models', *Proceedings of the National Academy of Sciences*, **108**, pp. 6306–11.

Ferguson, N.M., Laydon, D., Nedjati-Gilani, G., et al. (2020), 'Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand', Imperial College COVID-19 Response Team Report 9, March 16.

Forni, G. and Mantovani, A. (2021), 'COVID-19 vaccines: Where we stand and challenges ahead', *Cell Death & Differentiation*, **28**, pp. 626–39.

Galanti, M., Pei, S., Yamana, T.K., Angulo, F.J., Charos, A., Swerdlow, D.L. and Shaman, J. (2021), 'Social distancing remains key during vaccinations', *Science*, **371**, pp. 473–4.

Gavi: The Vaccine Alliance (2021), 'The COVAX facility: Interim distribution forecast latest as of February 3, 2021'.

Geoffard, P.-Y. and Philipson, T. (1997), 'Disease eradication: Private versus public vaccination', *The American Economic Review*, **87**, pp. 222–30.

Goodman-Bacon, A. (2018), 'Difference-in-differences with variation in treatment timing', Technical Report, National Bureau of Economic Research.

Goolsbee, A. and Syverson, C. (2021), 'Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020', *Journal of Public Economics*, **193**, p. 104311.

Hale, T., Angrist, N., Goldszmidt, R., et al. (2021), 'A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker)', *Nature Human Behaviour*, **6**, pp. 529–538.

Iboi, E.A., Ngonghala, C.N. and Gumel, A.B. (2020), 'Will an imperfect vaccine curtail the COVID-19 pandemic in the US?', Cambridge Working Papers in Economics 2097, Faculty of Economics, University of Cambridge. *Infectious Disease Modelling*, **5**, pp. 510–24.

Kremer, M. (1996), 'Integrating behavioral choice into epidemiological models of AIDS', *Quarterly Journal of Economics*, **111**, pp. 549–73.

Lee, J.Y. and Solon, G. (2011), 'The fragility of estimated effects of unilateral divorce laws on divorce rates', *The BE Journal of Economic Analysis & Policy*, **11**.

Makris, M. and Toxvaerd, F. (2020), 'Great expectations: Social distancing in anticipation of pharmaceutical innovations', Cambridge Working Papers in Economics 2097, Faculty of Economics, University of Cambridge.

McAdams, D. (2020), 'Economic epidemiology in the wake of COVID-19', *Economics*, **82120**, pp. 1–40.

McKee, M. and Rajan, S. (2021), 'What can we learn from Israel's rapid roll out of COVID 19 vaccination?', *Israel Journal of Health Policy Research*, **10**, p. 12.

Mecenas, P., Bastos, R.T.D.R.M., Vallinoto, A.C.R. and Normando, D. (2020), 'Effects of temperature and humidity on the spread of COVID-19: A systematic review', *PLoS one*, **15**, p. e0238339.

Pearl, J. (2013), 'Linear models: A useful "microscope" for causal analysis', *Journal of Causal Inference*, **1**, pp. 155–70.

Philipson, T. (2000), 'Economic epidemiology and infectious diseases', *Handbook of Health Economics*, **1**, pp. 1761–99.

Public Health England (2021a), 'Impact of COVID-19 vaccines on mortality in England: December 2020 to March 2021'.

Public Health England (2021b), 'Public Health England vaccine effectiveness report, March 2021'.

Rambachan, A. and Roth, J. (2019), 'An honest approach to parallel trends', Unpublished manuscript, Harvard University.

So, A.D. and Woo, J. (2020), 'Reserving coronavirus disease 2019 vaccines for global access: Cross sectional analysis', *British Medical Journal*, **371**.

Talamas, E. and Vohra, R. (2018), 'Go big or go home: a free and perfectly safe but only partially effective vaccine can make everyone worse off', Technical Report, Penn Institute for Economic Research, Department of Economics.

Toxvaerd, F. (2019), 'Rational disinhibition and externalities in prevention', *International Economic Review*, **60**, pp. 1737–55.

Toxvaerd, F. and Rowthorn, R. 2020, 'On the management of population immunity', Technical Report, Faculty of Economics & Cambridge-INET Institute, University of Cambridge.

Twohey, M., Collins, K. and Thomas, K. (2020), 'With First Dibs on Vaccines, Rich Countries Have 'Cleared the Shelves'", *The New York Times.*

Wright, A.L., Sonin, K., Driscoll, J. and Wilson, J. (2020), 'Poverty and economic dislocation reduce compliance with COVID-19 shelter-in-place protocols', *Journal of Economic Behavior & Organization*, **180**, pp. 544–54.

Yan, Y., Malik, A.A., Bayham, J., Fenichel, E.P., Couzens, C. and Omer, S.B. (2021), 'Measuring voluntary and policy-induced social distancing behavior during the COVID-19 pandemic', *Proceedings of the National Academy of Sciences*, **118**, p. e2008814118.

Zanin, L. and Marra, G. (2012), 'Rolling regression versus time-varying coefficient modelling: An empirical investigation of the Okun's law in some Euro area countries', *Bulletin of Economic Research*, **64**, pp. 91–108.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

Table A1. Descriptive statistics

	Mean	Standard deviation	Min	Max	*N*
Disease, vaccination and policy outcomes					
New cases	0.015	0.020	0	0.132	12,204
Deaths	0.046	0.048	0	0.227	12,076
OxCGRT Stringency index	0.618	0.157	0.120	0.935	12,204
Vaccinations per capita	0.020	0.068	0	1.102	12,204
Behavioural outcomes					
Retail	−25.081	20.707	−97	59	12,204
Residential	8.060	7.108	−14	55	12,204
Parks	−11.718	23.472	−89	134	12,204
Transit	−26.884	22.123	−98	59	12,204
Workplaces	−22.801	16.873	−94	74	12,204
Other variables					
GDP per capita	23.040	20.552	1.136	116.936	12,204
Temperature	0.110	0.114	−0.360	0.330	7,834
Humidity	0.734	0.147	0.110	1	7,834

Note: Behavioural outcomes from Google Mobility represent percentage changes relative to a baseline period in February 2020. GDP per capita is in thousands of 2019 U.S. dollars.

Table A2. Assessing effect of controlling for weather

	(1)	(2)	(3)	(4)
	Residential	**Residential**	**New cases**	**New cases**
Stringency index	0.022***	0.021***	−0.146***	−0.149***
	(3.48)	(3.37)	(−5.98)	(−6.07)
Vaccinations	−0.061**	−0.062**	−0.164*	−0.160*
	(−3.10)	(−3.18)	(−2.07)	(−2.02)
Temperature	−0.036***		0.132***	
	(−3.59)		(3.51)	
Humidity	0.020***		0.020	
	(5.69)		(1.53)	
N	7,760	7,760	7,752	7,752

Note: Estimates of residential visits and new cases, where both outcomes are in logs. In all models, the sample is restricted to observations with weather data available. Country-specific linear trends and complete sets of day and country effects are included in all models. Parentheses report *t*-statistics based on Driscoll–Kraay standard errors with up to 1 week arbitrary serial correlation. Asterisks denote statistical significance at the 0.05, 0.01 and 0.001 levels.

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

Table A3. Countries in estimation sample with selected summary statistics

	Vaccinations	New cases	Deaths	GDP
Afghanistan	0.0001	0.0003	5.7	1.8
Angola	0.0001	0.0002	1.3	5.8
Argentina	0.0096	0.0166	101.0	18.9
Australia	0.0006	0.0000	3.6	44.6
Austria	0.0281	0.0288	71.4	45.4
Bahamas	0.0004	0.0034	44.0	27.7
Bahrain	0.1000	0.0228	22.5	43.3
Bangladesh	0.0048	0.0007	4.7	3.5
Barbados	0.0319	0.0089	5.5	17.0
Belarus	0.0012	0.0165	16.7	17.2
Belgium	0.0268	0.0207	171.1	42.7
Belize	0.0015	0.0159	63.5	7.8
Bolivia	0.0013	0.0081	86.4	6.9
Bosnia and Herzegovina	0.0004	0.0195	127.0	11.7
Botswana	0.0000	0.0091	6.1	15.8
Brazil	0.0124	0.0221	100.8	14.1
Bulgaria	0.0097	0.0231	112.6	18.6
Cambodia	0.0019	0.0001	0.0	3.6
Canada	0.0203	0.0136	46.5	44.0
Cape Verde	0.0000	0.0095	22.5	6.2
Chile	0.0601	0.0157	93.3	22.8
Colombia	0.0014	0.0180	95.3	13.3
Costa Rica	0.0108	0.0142	45.9	15.5
Cote d'Ivoire	0.0000	0.0005	0.6	3.6
Croatia	0.0180	0.0349	99.1	22.7
Czech Republic	0.0246	0.0732	134.4	32.6
Denmark	0.0411	0.0227	28.6	46.7
Dominican Republic	0.0054	0.0085	24.3	14.6
Ecuador	0.0009	0.0058	82.6	10.6
Egypt	0.0002	0.0006	8.4	10.6
El Salvador	0.0007	0.0033	23.0	7.3
Estonia	0.0306	0.0490	26.3	29.5
Finland	0.0279	0.0072	10.8	40.6

(Continued)

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

Table A3. *Continued*

	Vaccinations	New cases	Deaths	GDP
France	0.0238	0.0266	103.4	38.6
Gabon	0.0000	0.0029	3.2	16.6
Georgia	0.0000	0.0419	65.7	9.7
Germany	0.0287	0.0178	53.8	45.2
Ghana	0.0013	0.0010	1.4	4.2
Greece	0.0289	0.0124	47.3	24.6
Guatemala	0.0002	0.0032	29.5	7.4
Honduras	0.0001	0.0063	35.0	4.5
Hungary	0.0343	0.0336	110.5	26.8
India	0.0040	0.0016	10.8	6.4
Indonesia	0.0035	0.0029	9.7	11.2
Iraq	0.0000	0.0052	32.1	15.7
Ireland	0.0325	0.0262	60.9	67.3
Israel	0.3911	0.0463	48.7	33.1
Italy	0.0284	0.0295	131.3	35.2
Jamaica	0.0004	0.0060	11.4	8.2
Japan	0.0002	0.0021	3.9	39.0
Jordan	0.0048	0.0290	38.6	8.3
Kazakhstan	0.0007	0.0051	15.2	24.1
Kenya	0.0000	0.0007	3.1	3.0
Kuwait	0.0222	0.0147	22.8	65.5
Kyrgyz Republic	0.0000	0.0028	21.1	3.4
Latvia	0.0083	0.0364	49.8	25.1
Lebanon	0.0018	0.0381	38.1	13.4
Lithuania	0.0340	0.0522	78.4	29.5
Luxembourg	0.0204	0.0434	81.6	94.3
Malaysia	0.0008	0.0070	2.1	26.8
Mali	0.0000	0.0002	1.4	2.0
Malta	0.0545	0.0353	53.5	36.5
Mauritius	0.0008	0.0002	0.8	20.3
Mexico	0.0061	0.0074	112.4	17.3
Moldova	0.0001	0.0243	79.3	5.2
Mongolia	0.0021	0.0010	0.1	11.8
Morocco	0.0286	0.0047	20.4	7.5
Mozambique	0.0001	0.0013	1.1	1.1

(Continued)

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

Table A3. *Continued*

	Vaccinations	New cases	Deaths	GDP
Myanmar	0.0043	0.0011	5.0	5.6
Namibia	0.0000	0.0089	11.3	9.5
Nepal	0.0086	0.0020	6.9	2.4
Netherlands	0.0232	0.0351	74.3	48.5
New Zealand	0.0005	0.0001	0.5	36.1
Nigeria	0.0000	0.0004	0.7	5.3
Norway	0.0312	0.0083	9.1	64.8
Oman	0.0063	0.0045	29.5	38.0
Pakistan	0.0002	0.0010	4.9	5.0
Panama	0.0084	0.0383	106.0	22.3
Papua New Guinea	0.0000	0.0002	0.1	3.8
Paraguay	0.0001	0.0129	35.2	8.8
Peru	0.0033	0.0121	122.6	12.2
Philippines	0.0001	0.0017	9.3	7.6
Poland	0.0311	0.0280	84.4	27.2
Portugal	0.0301	0.0488	98.9	27.9
Qatar	0.0313	0.0101	8.6	116.9
Romania	0.0295	0.0226	85.9	23.3
Russia	0.0150	0.0139	43.6	24.8
Rwanda	0.0023	0.0009	1.2	1.9
Saudi Arabia	0.0105	0.0007	17.9	49.0
Senegal	0.0008	0.0010	3.4	2.5
Singapore	0.0294	0.0003	0.5	85.5
Slovak Republic	0.0302	0.0382	70.3	30.2
Slovenia	0.0324	0.0588	138.5	31.4
South Africa	0.0003	0.0106	61.1	12.3
South Korea	0.0011	0.0011	2.2	35.9
Spain	0.0313	0.0313	120.4	34.3
Sri Lanka	0.0077	0.0027	1.3	11.7
Sweden	0.0298	0.0445	99.3	46.9
Switzerland	0.0233	0.0309	92.0	57.4
Thailand	0.0001	0.0003	0.1	16.3
Togo	0.0001	0.0006	0.9	1.4
Trinidad and Tobago	0.0002	0.0011	9.3	28.8

(Continued)

https://doi.org/10.1017/nie.2021.23 Published online by Cambridge University Press

Table A3. *Continued*

	Vaccinations	New cases	Deaths	GDP
Turkey	0.0347	0.0160	26.5	25.1
Uganda	0.0000	0.0005	0.6	1.7
Ukraine	0.0001	0.0188	48.1	7.9
United Kingdom	0.1072	0.0355	133.6	39.8
United States	0.0819	0.0460	120.7	54.2
Uruguay	0.0055	0.0162	9.7	20.6
Venezuela	0.0001	0.0014	4.0	16.7
Vietnam	0.0000	0.0000	0.0	6.2
Yugoslavia	0.0680	0.0536	49.1	14.0
Zimbabwe	0.0003	0.0015	5.6	1.9

Note: Table lists each country in the estimation sample along with vaccinations and mean new cases and total vaccinations per hundred capita per day, and total deaths to date per million capita. GDP is expressed in thousands of 2019 U.S. dollars.

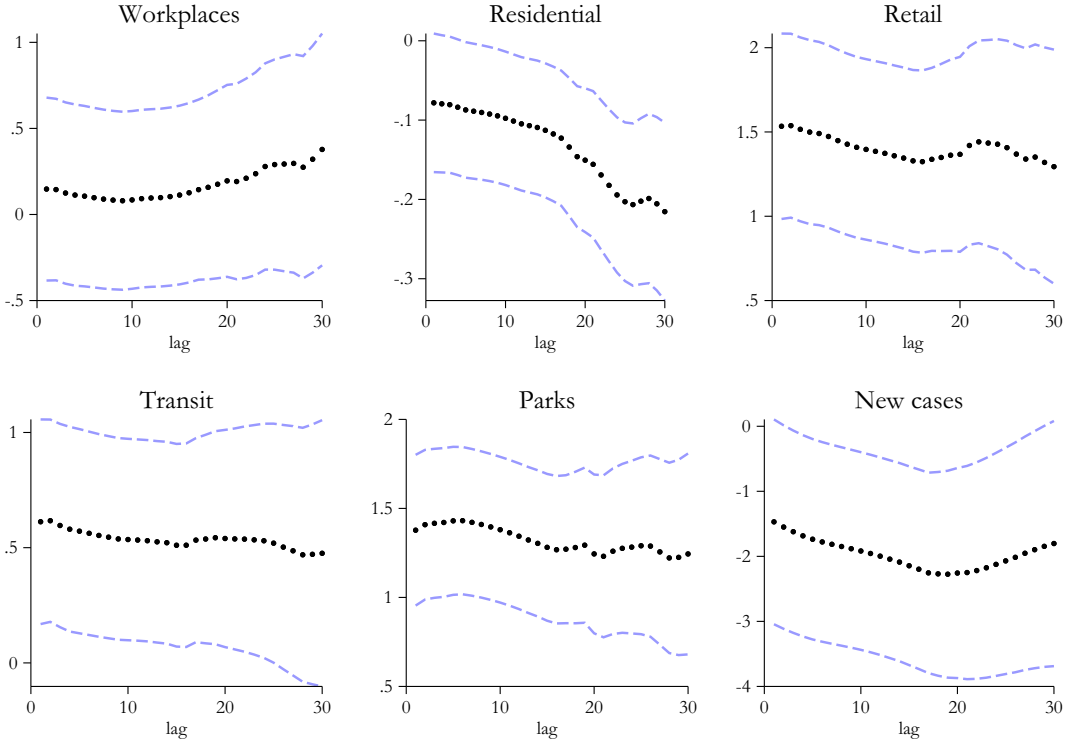

Figure A1. (Colour online) Sensitivity of estimates to lag length. Figure shows each model estimated in Table 2. Panel C, varying the lag of vaccinations per capita between 1 and 30. Note the scale of the *y*-axis varies across subfigures. The dashed blue lines are 95 per cent confidence intervals

Cite this article: Auld, M. C. and Toxvaerd, F. (2021), 'The great Covid-19 vaccine rollout: Behavioural and policy responses', *National Institute Economic Review*, **257**, pp. 14–35. https://doi.org/10.1017/nie.2021.23