Spatial Allocation of Scarce COVID-19 Vaccines^{*†‡}

François M. Castonguay,[§] Julie C. Blackwood,[¶] Emily Howerton,^{||}
Katriona Shea,^{||} Charles Sims,^{**} and James N. Sanchirico^{††}

July 11, 2021

Abstract

The COVID-19 Vaccines Global Access (COVAX) is an initiative led by the World Health Organization (WHO) and other partners that aims for an equitable access of COVID-19 vaccines. Despite a potential heterogeneous disease burden across space, countries receiving allotments of vaccines via COVAX may want to follow WHO's allocation rule and distribute vaccines to their jurisdictions based on the jurisdictions' relative population size. Utilizing economic–epidemiological modeling, we benchmark the performance of this *ad hoc* allocation rule by comparing it to the rule that minimizes the economic damages and expenditures over time, including a penalty cost representing the social costs of deviating from the *ad hoc* allocation. Under different levels of vaccine scarcity and different demographic characteristics, we consider scenarios where length of immunity and compliance to travel restrictions vary, and consider the robustness of the rules when assumptions regarding these factors are incorrect. The benefits from deviating are especially high when immunity is permanent, when there is compliance to travel restrictions, when the supply of vaccine is low, and when there is heterogeneity in demographic characteristics. Interestingly, a lack of compliance to travel restrictions pushes the optimal allocations of vaccine towards the *ad hoc* and improves the relative robustness of the *ad hoc* rule, as the mixing of the populations reduces the spatial heterogeneity in disease burden.

Key words: economic epidemiology, applied numerical methods, efficient disease intervention, allocation of scarce resources.

JEL Classification: C61, H12, H84, I18, Q54

***Authors' contributions:** F.M.C. and J.N.S. designed and performed the research, conceived and developed the analysis, F.M.C., J.C.B., E.H., and K.S. developed the epidemiological model, F.M.C. and E.H. parameterized the model, and F.M.C. ran the simulations. All gave feedback on the preliminary findings. F.M.C. wrote the first draft of the paper. All contributed to the final version.

[†]We thank for useful comments participants of the NatuRE Policy Lab at UC Davis (naturepolicy.ucdavis.edu), as well as participants of the internal seminars of the Department of Economics at *Université du Québec à Montréal* (UQAM), and those of the Montreal Environment and Resource Economics Workshop.

[‡]For the code used in this paper, see: <https://github.com/fmcastonguay/SpatialAllocationCOVID19>

[§]Department of Agricultural and Resource Economics, University of California, Davis, Davis, CA 95616, USA. Corresponding author. e-mail: fcastonguay@ucdavis.edu

[¶]Department of Mathematics and Statistics, Williams College, Williamstown, MA 01267, USA.

^{||}Department of Biology and Center for Infectious Disease Dynamics, Pennsylvania State University, University Park, PA 16802, USA.

^{**}Howard H. Baker Jr. Center for Public Policy and Department of Economics, University of Tennessee, Knoxville, Tennessee, TN 37996, USA.

^{††}Department of Environmental Science and Policy, University of California, Davis, Davis, CA 95616, USA and Resources for the Future, Washington DC 20036, USA.

1 Introduction

Now that several vaccines against coronavirus disease 2019 (COVID-19) have obtained the World Health Organization’s (WHO) emergency use listing, policymakers around the globe are deciding how to allocate the limited supplies within their boundaries. While the vaccination campaign of developed countries like the U.S. and the U.K. is well under way, developing countries in Asia and Africa face major vaccine supply issues. In anticipation of this problem, WHO and other partners created “COVID-19 Vaccines Global Access” (COVAX)—an initiative that aims for an equitable access of COVID-19 vaccines. The scientific literature has addressed the question of vaccine allocation, but most of it has focused on demographic considerations within one jurisdiction [1; 2; 3] or on a global scale [4; 5; 6; 7]. This prior work has made important contributions to the debate. A missing piece in the allocation question, however, is how to divide up limited quantities across jurisdictions (e.g. state/provinces, counties/regions) that might have different demographic and epidemiological characteristics. In the U.S., the National Academies of Sciences, Engineering, and Medicine (NASEM) recommends that vaccines are allocated to jurisdictions based on their relative population size “in the interest of speed and workability” [8] and WHO applies a similar principal with its COVAX program [9]. Countries whose vaccination campaign depends on COVAX provision could also apply a similar principle within their boundaries.

In this paper, we derive the optimal allocation and investigate its robustness to incorrect behavioral (i.e. compliance to nonpharmaceutical interventions) and epidemiological (i.e. the duration of immunity) assumptions. We also explore the economic and epidemiological trade-offs associated with a fixed *ad hoc* allocation rule similar to NASEM’s [8] and WHO’s [9] principle by comparing it to the optimal rule conditional on the level of scarcity of the vaccine. Furthermore, we investigate how various vaccine effectiveness—mimicking the wide range of effectiveness of the vaccines available for distribution (for the lower and upper end of the spectrum, see [10; 11] respectively) and the reduced effectiveness of vaccines when faced with emerging variants of concern [12]—affect the optimal allocation. The optimal rule we consider minimizes the economic costs from health-related damages, vaccine expenditures, and a workability cost imposed on the planner for deviating from the *ad hoc* rule.

In a world where two jurisdictions are identical in terms of population, the *ad hoc* rule would divide the limited supply equally between the jurisdictions. However, it is much more likely that two jurisdictions, even if equally sized, have heterogeneous levels of infections (e.g. in terms of cases) at the time a country’s central government receives an allotment of vaccines. Based on prior literature on spatial-dynamics of disease management, heterogeneity in infection levels may lead to significant deviations between the optimal spatial allocation and the *ad hoc* rule potentially leading to greater economic costs and worse public health outcomes (see [13] for example).

Mechanisms leading to heterogeneous infection include the timing of the outbreak, demographic characteristics of the population (e.g. age structure [14] and essential worker status [15]), and the implementation of and compliance with preventative nonpharmaceutical interventions; see [16; 17] for more details on how SARS-CoV-2 (i.e. the virus that

causes COVID-19) prevalence varies across space in the U.S. and see [18] for India, where the cumulative number of cases is highly heterogeneous across states. While compliance to preventive measures may seem independent from vaccine allocation, it affects the initial conditions (i.e. the conditions before the vaccine is allotted to the country) and the conditions under which the limited supplies will be allocated within the country. For example, compliance to shelter-in-place and travel restrictions results in little to no movement of the virus from one jurisdiction to another. When regions are non-interacting, Brandeau et al. [19] show for a general susceptible–infected–susceptible (SIS) model that the optimal allocation of resources depends on numerous intrinsic factors, including the size of the populations of each jurisdiction and the level of infection at the time of the vaccine allotment. When regions are interacting, Rowthorn et al. [20] show when there is no immunity (i.e. in an SIS model) that treatment should be preferentially directed towards the region that has the lower level of infection. While these results indicate that a fixed *ad hoc* rule is less cost-effective in an SIS model, whether compliance to travel restrictions makes the *ad hoc* rule relatively more cost-effective in the case of COVID-19 is an open question.

Our findings illustrate that the vaccines should be optimally allocated over time depending on: (i) if the jurisdiction has a lower or higher disease burden at the time of the vaccine allotment, (ii) if immunity is permanent (see Zhou et al. [21]) or temporary (Gersovitz and Hammer [22] already pointed out that the optimal allocation is conditional on the duration of immunity), (iii) whether there is compliance to travel restrictions or not, (iv) the amount of vaccine available, and (v) the average demographic characteristics of the population (i.e. age structure and essential worker status). We proxy variability in demographics by assuming that the population of one jurisdiction has a higher case-fatality ratio (e.g. an older population [14]) or a higher contact rate (e.g. a population with more essential workers [15]) than the other. We find that the benefits of deviating from the *ad hoc* rule—which prioritizes equity of distribution—are especially high when immunity is permanent, when there is compliance to travel restrictions, when the vaccine supply is low, and when there is heterogeneity in demographic characteristics. Allocating a vaccine based on an *ad hoc* allocation rule generally leads to an over-utilization in jurisdictions where disease prevalence is higher, an under-utilization in jurisdictions where disease prevalence is lower, and overall a higher number of cumulative cases. Whether these inefficiencies outweigh the “speed and workability” [8] inherent in *ad hoc* rules is an important question for policymakers. Our research can aid in that discussion by illuminating the trade-offs involved in such complex epidemiological, economic, and social decisions by providing optimal benchmarks from which to compare *ad hoc* rules.

While the optimal allocation is conditional on a number of factors mentioned above, the science remains unresolved on the duration of immunity to SARS-CoV-2, and it is difficult to anticipate and subsequently estimate the extent to which populations in different jurisdictions comply with the travel restrictions. On the other hand, the *ad hoc* allocations have the advantage of being based on easily observable factors (e.g. a jurisdiction’s population size). To gain insights into the robustness of optimal and *ad hoc* policies in the presence of such uncertainties, we investigate

the economic and public health consequences that could occur if we design an optimal policy or evaluate the performance of *ad hoc* rules under a set of assumptions on immunity and compliance that turn out to be incorrect.

We make a number of contributions to the literature. First, we develop an economic–epidemiological model and solve for the optimal allocation of vaccines over time to minimize the economic costs from health-related damages, vaccine expenditures, and a workability cost imposed on the planner for deviating from the *ad hoc* rule. Prior literature considering the trade-offs involved with *ad hoc* rules does not consider that deviating from them entails potential workability costs (see, for example, [23]). Second, we consider how vaccine allocations are influenced by compliance with preventative nonpharmaceutical interventions (i.e. travel restrictions). Third, we demonstrate how vaccine allocations are dependent on various demographics (i.e. age structure and essential worker status). Fourth, we show that, in general, optimal rules are robust to incorrect assumptions about the duration of immunity but differences in public health outcomes (cumulative cases) appear when compliance to a travel restrictions is assumed when in fact there is not compliance; it is, however, much preferable from a public health outcome perspective to comply with travel restrictions.

The paper is divided as follow. In Section 2, we detail the different types of interventions, we present the components of the economic-epidemiological model, and detail the technique used to analyse the allocation question. Section 3 presents the results while Section 4 concludes the paper.

2 Material and Methods

We develop an economic–epidemiological model to describe the dynamics of SARS-CoV-2. The model captures a situation where a central planning agency (e.g. the central government) must decide when and how much of the scarce vaccines to allocate to two jurisdictions where disease burden is heterogeneous at the moment it receives an allotment of vaccine. We assume that the objective of the central planner is to minimize costs across both jurisdictions, including damages associated with the morbidity and deaths of infected individuals, the expenditures related to the pharmaceutical intervention, and a penalty cost mimicking the increased workability costs incurred for any deviation from the *ad hoc* allocation. The dynamics of SARS-CoV-2 are modeled using an SEIR epidemiological model, which tracks the change over time of the susceptible (S), exposed (E), infected (I), and recovered (R) populations for two separate jurisdictions (see Appendix A for more details on the calibration of the model). We note that while we generally talk about these jurisdictions as being two different states, they can very well represent any two sub-national jurisdictions like provinces or territories, and even counties or regions within one sub-national jurisdiction.

2.1 Modelling Different Types of Intervention

There are two different types of interventions we consider: travel restrictions and vaccines. We assume that travel restrictions affect both jurisdictions simultaneously (e.g. by an order from the central government), and that the populations either comply perfectly or imperfectly to the travel restrictions (for examples of optimal lockdown policies see, e.g., [24; 25]). When compliance is perfect, individuals in different jurisdictions do not interact with each other and thus susceptible individuals can only get infected by being in contact with some infected individual in their own jurisdiction. When compliance is imperfect, susceptible individuals from one jurisdiction can also travel to the other jurisdiction where they can be in contact with infected individuals, or infected individuals from one jurisdiction can travel to the other jurisdiction and infect susceptible individuals there; this discrete shift in the number of contacts effectively increases the transmissibility of the virus (see Appendix A more details).

We assume that the analysis starts when the central planning agency has received an allotment of vaccines and will keep receiving a continuous allotment of vaccines. For simplicity, the amount of available vaccine is assumed to be exogenous to the model and fixed over time, which is likely given the short time frames we consider in the paper. However, we consider different levels of vaccine allotments, or capacity, to investigate how different levels of vaccine scarcity may affect their optimal allocation. In our model, vaccines reduce the pool of susceptible individuals by providing them with immunity from the virus, as early evidence suggests that vaccines could be transmission blocking in addition to preventing severe disease [26].

2.2 Model of Disease Transmission

We use a frequency-dependent [27] susceptible–exposed–infected–recovered (SEIR) model that describes the dynamics of COVID-19 in two separate jurisdictions $i = 1, 2$ (e.g. states/provinces or counties/regions); each jurisdiction contains a population of N_i individuals that is either susceptible, exposed, infected, or recovered (see Figure 1). We also consider scenarios where immunity is temporary (i.e. lasts 6 months, for more details see [28]), thus also using an SEIR–Susceptible (SEIRS) model (for COVID-19 applications see, e.g., [29; 30; 31; 32; 33; 34; 35; 36; 37; 38; 39]). In such scenarios, the R_i recovered individuals are immune for a mean period of $\frac{1}{\omega}$ months.

In each jurisdiction i , the S_i susceptible individuals are in contact with the I_i infected individuals of their own jurisdiction at a rate of β_{ii} and are in contact with the I_j infected individuals of the other jurisdiction at a rate of β_{ij} . We assume $\beta_{ij} = 0$ (i.e. no mixing between jurisdictions) when there is perfect compliance to travel restrictions, and $\beta_{ij} > 0$ if not. To highlight the role of travel restriction compliance and initial disease burden, we initially assume that the contact rate is identical across jurisdictions, meaning that $\beta_{11} = \beta_{22} = \beta_{ii}$ and $\beta_{12} = \beta_{21} = \beta_{ij}$ (in Section 3.2 we relax this assumption and investigate the optimal allocation when there is heterogeneity in the contact rate). We assume there is no permanent migration of individuals from one jurisdiction to another (see for instance [40] and see [41] for an example applied to COVID-19) in the sense that individuals who do not comply with

travel restrictions do not permanently move to the other state, but instead travel to it temporarily. An implication is that we are assuming that the two jurisdictions are close enough for such travel and mixing to be economically feasible.

We model the control variables for vaccines as non-proportional controls, i.e. available in a constant amount each month [3; 20; 42; 43]. The change in susceptible individuals is

$$\dot{S}_i = \omega R_i - \beta_{ii} S_i \frac{I_i}{N_i} - \beta_{ij} S_i \frac{I_j}{N_j} - q_V u_{V_i} \quad (1)$$

where u_{V_i} represents the number of individuals being treated via vaccine in a given time period (i.e. a month) in Jurisdiction i , and q_V represents the effectiveness of the vaccine (note that we look at the low end of vaccine effectiveness to be conservative; see Section 3.4 for more details). We note that our model does not distinguish between individuals whose vaccine has failed and those who have not been vaccinated at all. As such, individuals with vaccine failure can be re-vaccinated in subsequent months.

After being infected, susceptible individuals transition into the exposed class E_i where the disease remains latent for a mean period of time of $\frac{1}{\sigma}$, before the onset of infectiousness. The change in the number of exposed individuals is

$$\dot{E}_i = \beta_{ii} S_i \frac{I_i}{N_i} + \beta_{ij} S_i \frac{I_j}{N_j} - \sigma E_i. \quad (2)$$

Exposed individuals eventually become infectious for a mean period of time of $\frac{1}{\gamma + \varphi_i}$ and in turn can infect susceptible individuals. Infected individuals either recover naturally from the disease at a rate of γ or die from complications related to infection at a disease induced mortality rate of φ_i . In our base case we assume identical disease induced mortality rates across jurisdictions, i.e. $\varphi_1 = \varphi_2 = \varphi$ but investigate the optimal allocation when $\varphi_1 \neq \varphi_2$ in Section 3.2. The growth of the infected individuals is

$$\dot{I}_i = \sigma E_i - \gamma I_i - \varphi_i I_i. \quad (3)$$

The recovered population R_i includes individuals that recover naturally from the disease at a rate of γ and the individuals that are successfully vaccinated every month ($q_V u_{V_i}$); if immunity is temporary ($\omega > 0$), a fraction of the recovered will leave this compartment. Our model does not distinguish between vaccine-acquired immunity and naturally-acquired immunity. The number of recovered individuals in Jurisdiction i thus changes according to

$$\dot{R}_i = \gamma I_i + q_V u_{V_i} - \omega R_i. \quad (4)$$

At any instant in time, we have that $N_i = S_i + E_i + I_i + R_i$, which in turn implies that the growth of the population over time is

$$\dot{N}_i = -\varphi_i I_i. \quad (5)$$

In keeping with much of the previous economic epidemiology literature [22] as well as recent applications to COVID-19 (see for example [44]), we have omitted natural births and non-COVID-related deaths due to the short time frame of our model (4 months) and assume reductions in international travel [45] effectively lead to a closed population (i.e. there is no exogenous importation of infected individuals). See Appendix A for more details about the parameterization of the epidemiological model.

Figure 1: This schematic shows the model interventions and disease transmission pathways for our model of COVID-19. The full lines represent the transition between, or out of, compartments while the dotted lines represent contact between susceptible and infected individuals. Black lines represent situations that do not vary, while yellow lines represent key factors that we vary in our model to see how they impact our results. The green line represents the vaccines and the red line represents mortality.

2.3 Modelling *Ad Hoc* Allocations

We model an *ad hoc* allocation rule that favors “speed and workability” [8]. We follow the NASEM [8] and WHO [9] principle and impose that the allocation is based on relative population sizes. Specifically, the rule for Jurisdiction

i is that

$$u_{V_i} \leq \left(\frac{N_i}{N_1 + N_2} \right) \bar{u}_V \quad (6)$$

where $\bar{u}_V$ is the limited amount of vaccine allotted to the central government. When the population sizes are the same, the *ad hoc* rule will divide equally the limited doses to the two jurisdictions.

In the *ad hoc* scenarios, we model the allocation rule as an inequality because towards the end of the horizon after periods of vaccinations, the level of susceptible in the population may be such that the limited supply of vaccines is not an issue. Other *ad hoc* rules are possible, such as, allocate all to the largest or smallest population [23], but we concentrate on the one currently being advocated for by NASEM [8] and WHO [9].

2.4 Model of Economic Costs

The model of economic costs include damages related to morbidity and deaths, costs spent on the vaccines, and the workability cost described above that is incurred for any deviation from the *ad hoc* allocation rule. Damages represent consequences related to a temporary disability associated with severe or critical symptoms, and loss of life in the worst cases. The damages are assumed to be linear and additively separable across jurisdictions, meaning that they are identical across individuals and across jurisdictions. The marginal value of damages (i.e. the damages associated with the death of one individual) is assumed to be constant over time and given by the value of a statistical life (VSL) that the U.S. Environmental Protection Agency [46] uses (see Appendix A for more details on the parameterization). Damages incurred from a temporary disability associated with severe or critical symptoms can be compared to deaths via some disability weight w ; given we found no published disability values associated with COVID-19, we follow the literature (see for instance [47]), and use the disability value associated with lower respiratory tract infections. The damage function for Jurisdiction i is

$$c_i(I_i) = (w + \varphi_i)cI_i \quad (7)$$

where c is the damage parameter associated infectious individuals (i.e. the VSL).

We model a scenario where the central planner is focused on the allocation of vaccines where the costs for its development have already been incurred. This implies that vaccine development costs have already been utilized (in technical terms we say that the costs are sunk) and therefore do not affect the decision of the central planning agency. We model the vaccination cost as linear, where the cost parameter represents the cost of vaccinating one individual. The vaccine cost function is denoted $c_{V_i}(u_{V_i})$, with $i = 1, 2$. We assume that the vaccination cost is

additively separable across jurisdictions such that we denote the cost of vaccinating u_{V_i} individuals as

$$c_{V_i}(u_{V_i}) = c_V u_{V_i} \text{ for } i = 1, 2, \quad (8)$$

where c_V represents the cost of treating one individual via vaccine. Note that calibration of the cost parameter is based on current vaccine prices (see Appendix A for more details about the parameterization of the economic model), but it could also represent the cost the central planner needs to incur to pay to administer the vaccine.

We assume that the central planning agency incurs a workability cost representing the social (transaction) costs of deviating from the *ad hoc* allocation rule (for another application of this concept, see [48]). The workability cost function is:

$$c_A(u_{V_1}, u_{V_2}, N_1, N_2) = c_A \left(\left(\frac{N_2}{N_1 + N_2} \right) u_{V_1} - \left(\frac{N_1}{N_1 + N_2} \right) u_{V_2} \right)^2 \quad (9)$$

where c_A is the parameter associated with the workability cost. When the gains from deviating from the *ad hoc* allocations (i.e. a reduction in damages in one jurisdiction) outweigh the costs (i.e. an increase in damages in the other jurisdiction and the increased workability costs incurred), the central planning agency will prioritize this allocation as it will lead to lower total costs. By imposing the *ad hoc* rule *ex ante*, the decision-maker is essentially assuming that this workability cost is infinite. Everything else being equal, we expect that the presence of the workability cost will push the optimal allocation towards the *ad hoc* rules (see Figure A18 for a sensitivity analysis of our results to the workability cost parameter). Therefore, when we do find deviations, we need to consider that these include this workability cost and if workability costs smaller, then the deviations and trade-offs would be greater.

2.5 Planner's Objective

In optimal control theory, the best, or optimal, path of the control variables (here the allocation of the limited supply of vaccines) is conditional on the objective of the central planning agency. We assume that the objective is to minimize the economic damages and the costs of the pharmaceutical intervention across jurisdictions over time, rather than a solely epidemiological objective (see for instance [20]). The objective function is the net present value of damages, expenditures related to vaccination, and the workability cost over an exogenously determined planning horizon (4 months). Specifically, the planner's objective is:

$$\min_{u_{V_1}, u_{V_2}} \int_0^T e^{-rt} \left\{ c_1(I_1) + c_2(I_2) + c_{V_1}(u_{V_1}) + c_{V_2}(u_{V_2}) + c_A(u_{V_1}, u_{V_2}, N_1, N_2) \right\} dt \quad (10)$$

where r is monthly discount rate. The planner solves equation (10) over a fixed time interval, T , subject to equations (1), (2), (3), (4), (5), along with constraints on availability of vaccines ($u_{V_1} + u_{V_2} \leq \bar{u}_V$), non-negativity conditions,

physical constraints on vaccines, initial disease burdens in each jurisdictions, and free endpoints (see discussion on terminal conditions in the next section). In the *ad hoc* scenarios, we also impose equation (6).

2.6 Initial and Terminal Conditions

The disease burden in each jurisdiction at the beginning of the time horizon (i.e. in $t = 0$ when the allotment of vaccine is received for the first time) is calibrated using the epidemiological model (equations (1), (2), (3), (4), and (5)). At the beginning of the outbreak, we assume that, in each jurisdiction, there is one exposed individual in an otherwise entirely susceptible population of 10 million individuals, and that populations of the different jurisdictions comply with the travel restrictions. The only difference between the two jurisdictions is that the outbreak started one week earlier in State 2. We simulate the outbreak for approximately nine months to yield the initial conditions; see Appendix B for more details. In Section 3.2 when we consider heterogeneity in demographic characteristics (varying case-fatality ratio and contact rate), we modify the initial conditions accordingly assuming an identical timing in the outbreak of the disease.

We impose no conditions on the number of susceptible, exposed, infected, and recovered individuals at the end of the planning horizon; in technical terms, we say that the state variables are free (see Appendix B for more details). Under our free endpoint conditions, there is a transversality condition (i.e. a necessary condition for the vaccine allocation to be optimal) for each state variable that requires the product of the state variable (S_i, E_i, I_i, R_i or N_i) and its corresponding costate variable (i.e. the shadow value, or cost, associated with the state variable) is equal to zero. Hence, at the end of the time horizon, either the state variable equals zero, the shadow value associated with the state variable equals zero, or both. In any case, allowing state variables to be free guarantees that the terminal levels of the state variables are optimally determined. Another possible assumption could be that over a fixed interval we find the optimal policy such that at the end of the horizon there is a given percent reduction in infected or susceptible individuals. Our approach nests this more restricted scenario.

3 Results

To examine the optimal allocations of vaccine over time, we numerically solve the optimal control problem across three different scenarios: no controls, optimal vaccine allocation, and *ad hoc* vaccine allocation. We investigate how to allocate vaccines by mapping out the different allocation rules for different immunity–travel restrictions–capacity scenarios. Any deviation from the *ad hoc* allocation rule is optimal despite incurring the workability cost. As the workability cost parameter c_A goes to zero, the problem becomes linear in the controls where the optimal allocations in linear problems follow singular solutions. We use pseudospectral collocation to solve for the optimal dynamics of vaccine and infection over time, which converts the continuous time optimal control problem into a constrained

non-linear programming problem solving for the coefficients of the approximating polynomials at the collocation nodes (see [49; 50] for other applications, and see Appendix B for more details on this technique).

We present the results for our preferred specification of the parameters (i.e. following what was estimated in the literature; see details in Appendix A) and for the case where immunity is permanent and the case where immunity is temporary. We detail the optimal deviation based on whether the populations of the different jurisdictions are compliant to travel restrictions or not, and for different levels of capacity constraints. The total available quantity of vaccine in a given time period (i.e. a month; $\bar{u}_V$) is based on a certain percentage (5%, 10%, or 15%) of the total population size. We focus our analysis on the period of time when the scarcity of the vaccine constraint is binding, as once the constraint relaxes the allocation question becomes moot.

3.1 Base Case: Homogeneous Demographic Characteristics

Compliance to travel restrictions impacts the optimal allocation of vaccines, regardless of whether immunity is temporary or permanent and regardless of the amount of vaccine available. Noncompliance to travel restrictions reduces both the oscillation (i.e. back-and-forth movement of resources between jurisdictions) of the optimal allocation and the amplitude of the deviations from the *ad hoc* rule (see Figure 2 for when immunity is permanent, and see Figure A1 for when immunity is temporary). Because noncompliance to travel restrictions decreases the structural heterogeneity in the system, the optimal allocation of vaccine converges towards the *ad hoc* allocation when populations mix with each other. This result clearly demonstrates how the performance of the allocation rule is dependent on how citizens in the jurisdictions comply with travel restrictions.

Figure 2: **Vaccine allocation with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether there is compliance to travel restrictions (panels A and C) or not (panels B and D) for the case where the vaccine capacity constraint is 10% and immunity is permanent. Note the changing *y*-axis in panels C and D in order to better highlight the infection levels.

Noncompliance to travel restrictions leads to the initially less infected state being favored by the optimal allocation for low levels of vaccine capacity (e.g. 5% capacity; see Figure 3 Panel A for when immunity is permanent and Figure A4 Panel A for when immunity is temporary). On the other hand, the more infected state will be prioritized at the beginning of the time horizon for a very short period of time when vaccine capacity is larger (e.g. 10% or 15% capacity; see Figure 3 panels B and C for when immunity is permanent and Figure A4 panels B and C for when immunity is temporary). More generally, regardless of whether or not populations are compliant with travel restrictions, and regardless of whether immunity is temporary or permanent, a higher vaccine capacity implies that relatively more of the supply should be given to the more infected state at the beginning of the time horizon (see figures 3 and A2 for the case where immunity is permanent; see figures A3 and A4 for the case where immunity is temporary).

Figure 3: **Vaccine allocations under different levels of scarcity without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether capacity is 5% (panels A and D), 10% (panels B and E), or 15% (panels C and F), for the case where immunity is permanent and there is no compliance to travel restrictions.

Interestingly, temporary immunity has a different effect on the optimal vaccine allocation depending on whether or not populations are compliant to travel restrictions. When populations comply with travel restrictions, temporary immunity increases the oscillation of the optimal allocation because benefits from vaccination are only temporary, and since the population gradually loses its immunity, it forces more back-and-forth movement of resources between jurisdictions (see Figure A5). When populations do not comply with travel restrictions, temporary immunity reduces the amplitude of the deviations from the *ad hoc* rule because it further dampens the structural heterogeneity in the system, since the infection and recovery level of both jurisdictions will eventually reach the same positive steady-state level (recall the only heterogeneity in the system is the initial disease burden in the base case).

While the optimal allocation of vaccine is unequal from a resource allocation perspective, it equalizes the current infection levels across jurisdictions (Figure 2 Panel C), that is, the optimal allocation seems to be prioritizing equity of outcomes as opposed to the *ad hoc* allocation, which prioritizes equity of distribution. As the vaccine capacity increases, however, the *ad hoc* allocation rule performs better and in turn the amplitude of the optimal deviation decreases (see Figure A2 panels A, B and C, or Figure A3 panels A, B, and C). These optimal cost-minimizing deviations that lead to equal current infection levels across jurisdictions towards the end of the time horizon imply

that the optimal cumulative number of cases is more unequal than in the *ad hoc* allocation (Figure 4). Hence, the optimal allocation makes the current infection level more equal, while the *ad hoc* allocation makes cumulative infection more equal. In fact, in all scenarios considered, the optimal allocation will lead to lower cumulative damages in the less infected jurisdiction but higher cumulative damages in the most infected jurisdiction (see Figure 4 with vaccine capacity of 10%, and see figures A7 and A8 with vaccine capacity of 5% and 15%, respectively).

Figure 4: **Epidemiological outcomes under different scenarios with vaccines.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-vaccine case for different immunity–travel restrictions (TR) scenarios and for when vaccine capacity is 10%.

3.2 Heterogeneous Demographic Characteristics

In our main analyses, we introduced heterogeneity in infection across jurisdictions by assuming a different timing in the outbreak of the virus. While this leads to heterogeneity in the disease burden, there are other mechanisms that could lead to similar differences in the disease burden at the time the allotment of vaccine is received. For instance, in jurisdictions that have an older population on average, we expect SARS-Cov-2 to have a higher case-fatality ratio [14] which would effectively decrease the transmissibility of the virus (see Appendix A more details). If we assume that one jurisdiction has a higher case-fatality ratio and start the initial outbreak at the same time, then we find the population with the highest case-fatality ratio also has the largest population of susceptible individuals at the time the vaccine is administered. The heterogeneity in initial disease burden stems from the lower transmissibility

in the jurisdiction with the higher case fatality ratio. The heterogeneity in case fatality ratio therefore not only leads to heterogeneity in the initial disease burden but it also implies that the benefits of vaccination are no longer homogeneous across jurisdictions. These differences lead the optimal allocation to favor even more the least burdened jurisdiction, which is also the most vulnerable (aka more older individuals) of the two populations. Overall, introducing heterogeneity in case fatality ratio strengthens our main set of results (see Figure A9 for when immunity is permanent and see Figure A10 for when immunity lasts 6 months).

Another source of heterogeneity in infection could stem from one jurisdiction having more essential workers than the other (for more details on how the risk of infection is occupation-dependent, see [15]). In our model, we can capture this by considering spatial heterogeneity in the contact rate, where a higher contact rate proxies for more essential workers. This in turn leads to a higher initial disease burden in the jurisdiction with higher contact rate. We find that priority is given to the state with a higher contact rate (aka more essential workers) in almost all cases. As the state with a higher contact rate gets vaccinated, we eventually shift priority to the state with the lowest contact rate because either the number of cases starts decreasing in the jurisdiction that has the higher contact rate, or the low-contact state eventually reaches a point where its infection level becomes higher than the high-contact state (see Figure A11 for when immunity permanent and see Figure A12 for when immunity lasts 6 months). In the case when immunity lasts 6 months and there is compliance to travel restrictions (Figure A12, Panel A), we find that priority is given to the low-contact jurisdiction, which contradicts many notions of fairness associated with vaccine allocation. With the gains of vaccination temporary and no movement of people, it turns out that the greatest return per vaccination is in the place where we can best avoid future cases (low contact rate jurisdiction). This prioritization is only fleeting however and there is more back-and-forth movement of resources between jurisdictions in this case, even though the workability cost is being incurred each time.

3.3 Robustness of Spatial Allocations

There is significant uncertainty associated with the duration of immunity (i.e. if it is permanent or temporary) and to what extent populations comply with travel restrictions. One argument for the *ad hoc* allocation is that uncertainty in these parameters makes the optimal allocation impossible to achieve. This uncertainty is not yet resolved and public health officials have to choose vaccine allocations based on potentially incorrect assumptions. We compare the robustness of the optimal spatial allocation to the *ad hoc* allocation. By definition, the optimal allocation minimizes the net present value of the health-related damages and total expenditures (including vaccine expenditures and the workability cost incurred because of the deviations from the *ad hoc* allocation), and thus cannot do worse on this dimension than the *ad hoc* allocation. We measure robustness by first inserting the optimal solution under one set of assumptions into the disease dynamics under another set and compute the changes in total expenditures (i.e. the pharmaceutical intervention and the workability cost) and public health outcomes (cumulative cases) over time. We then calculate the distance of these changes in percentage terms to the optimal

solution derived under the “correct” assumptions (represented by the point (0,0) in Figure 5). For example, suppose immunity is permanent and there is perfect compliance to a travel restriction. We derive the optimal policy under these assumptions and use it to measure the robustness of the optimal policies that are derived under assumptions that immunity is temporary and/or there is noncompliance. The *ad hoc* policies being based on observable factors are then compared to the incorrectly applied optimal policies. We illustrate the case for 10% scarcity and include other scarcity cases in Appendix C.

When demographic characteristics are homogeneous across jurisdictions, we find overall that immunity length has a lesser impact on both economic and epidemiological outcomes than compliance to travel restrictions (compare the distance from the origin between the plusses and the stars in Figure 5). There are more nuanced trade-offs, however (e.g. compare position of the stars across the panels in Figure 5). Across the economic dimension (expenditures), for example, we find that assuming compliance when in fact there is very little leads to greater expenditures. Recall by design, the *ad hoc* allocations have lower expenditures than the optimal policies because the central planner is not incurring the workability costs from deviating off of the allocation. At the same time, greater cumulative cases result when the opposite holds, that is, assuming no compliance when in fact there is compliance. We also see that in some instances that the combined effect of incorrectly assuming the wrong immunity and compliance can offset some deviations (e.g. see Figure 5 Panel C) while in other cases the results are dominated by non-compliance. Finally, when there is compliance to travel restrictions the *ad hoc* allocation performs worse than any of the optimal allocations, while the *ad hoc* allocation performs relatively well when there is no compliance to travel restrictions. Varying the level of scarcity does not change the qualitative nature of results (see figures A13 and A14 for when vaccine capacity is 5% and 15% respectively), except for one anomaly where the *ad hoc* does not always perform worse under assumptions on compliance to travel restrictions (Figure A14).

Figure 5: **Robustness of epidemiological and economic outcomes under different scenarios with vaccines.** Percentage change in expenditures (y -axis) and percentage change in cumulative cases (x -axis) from the optimal allocation for different immunity-travel restrictions scenarios and for when vaccine capacity is 10%. The x -axis represent small percentage changes but when scaled up to population level effects translate into significant differences in public health outcomes.

We also investigate the robustness of the optimal allocations when the demographic characteristics are heterogeneous across jurisdictions. When jurisdictions have a different case-fatality ratio, the *ad hoc* allocation performs better than the optimal allocations when considering cases as the main health outcome (Figure A15). However, this approach is misleading because when case-fatality ratios are heterogeneous across jurisdictions, the cumulative aggregate number of cases (all jurisdictions together) is a poor outcome measure as a case in one place is not the same as a case in another jurisdiction. In this setting, the disease burden and cumulative damages give a more accurate depiction of the situation. In fact, while the *ad hoc* allocation outperforms the optimal allocations in terms of cumulative cases, it performs considerably more poorly when considering cumulative damages. We generally find that the optimal allocations outperform the *ad hoc* allocation in all scenarios considered (Figure A16). When jurisdictions have a different contact rate, the same pattern as in Figure 5 holds in the sense that when there is compliance to travel restrictions, the optimal allocations outperform the *ad hoc* allocation, while the *ad hoc* generally performs better than the optimal allocations when there is noncompliance to travel restrictions (Figure A17).

3.4 Sensitivity Analyses

The previous section considers the robustness of optimal allocations to incorrect assumptions about parameters (e.g. assuming permanent immunity while in fact it is temporary). Public health officials will also want to know how much optimal allocations change when parameters change (e.g. because vaccine effectiveness is lower against a new strain of the virus). We address those questions in this section. While both sets of analyses address parameter uncertainty, you can consider in this section that the uncertainty is resolved before the public health officials have to make the vaccine allocation, while in the previous section the uncertainty was not resolved and public health officials had to choose allocations based on potentially incorrect assumptions.

Two key parameters in our analysis are the scale of the workability cost (c_A in Equation (9)) and the level of vaccine effectiveness (see Appendix C for more details). While imposing the *ad hoc* rule *ex ante* implicitly means that the cost of deviating from the *ad hoc* allocation is infinite, in practice it is likely finite but hard to quantify, as it depends on logistical, political, and cultural factors. We investigate the sensitivity of our results by solving for optimal vaccine allocation over a range of values. We find greater deviations off of the *ad hoc* at lower workability costs resulting in greater differences in cumulative cases, and smaller deviations as the workability cost parameter increases (Figure A18 panels A, B, C, and D). Specifically, we find that when the cost is in the neighborhood of the VSL (c in Equation (7) and Figure A18 black line represents the VSL), that the planner no longer deviates from the *ad hoc*.

The base case parameter for vaccine effectiveness we utilized in the paper is based on estimates of the influenza vaccine [51] (see Appendix A for more details) and represents a conservative estimate similar to the lower efficacy vaccines listed for emergency use listing by WHO (i.e. the Sinopharm COVID-19 vaccine, see e.g. [10]). Evidence from other COVID-19 vaccines (e.g. Pfizer/BioNTech and Moderna vaccines) suggest that effectiveness could be considerably higher than our base case scenario [11], but preliminary evidence suggests that vaccine efficacy is lower against new variants (e.g. the Delta variant, aka lineage B.1.617.2; see e.g. [12] and see [52] for a detailed report). We find that the more effective a vaccine is, the more a central planner would want to deviate from the *ad hoc* allocation (in blue; Figure A19 panels A, B, C, and D). As a result of this greater deviation, we see a larger difference in terms of the reduction in cumulative cases (in red; Figure A19 panels A, B, C, and D).

4 Conclusion

Recent studies have discussed how a vaccine against the coronavirus disease (COVID-19) should be allocated within a geographical area (see for instance [1; 2; 3]) and on a global scale (see for instance [4; 5; 6; 7]). Building off the spatial-dynamic literature in epidemiology, we contribute to this body of work by addressing the question of distributing a scarce allotment of COVID-19 vaccine across smaller geographic areas, such as counties or states,

and by showing how optimal allocations perform when vaccine allocation must be determined before uncertainty about key behavioral (i.e. compliance to travel restrictions) and epidemiological (i.e. the length of immunity to the disease) factors are resolved. Countries receiving allotments of vaccines via COVAX—a World Health Organization (WHO) led initiative aimed to provide an equitable access to COVID-19 vaccines—could follow NASEM’s [8] and WHO’s [9] principle and allocate a COVID-19 vaccine to jurisdictions within their borders based on the jurisdictions’ population size. While this approach prioritizes an important principle—equity of distribution—we show in this paper that there are potential economic and public health benefits of deviating from such an allocation rule. The optimal allocation rule we derive is more closely aligned with another important principle—equity of outcomes—and this paper provides policymakers with information on the trade-offs involved with different allocation rules. We also highlight under which behavioral and epidemiological conditions the optimal allocation performs relatively better and when the *ad hoc* allocation is a good approximation of the optimal allocation. There are many factors that come into play in these allocation decisions and the methodology proposed here provides a way to benchmark these rules to illustrate the trade-offs. Other methodologies that do not solve for the optimal policies are left to benchmark one set of *ad hoc* rules against another, where the set of possible *ad hoc* rules is infinite.

We considered several different scenarios where the length of immunity, the compliance to travel restrictions, the size of the vaccine allotment, and the demographics across jurisdictions are varied. In most of these scenarios, we find that priority should be given to jurisdictions that initially have lower disease burden. The intuition behind this result—already put forward by Rowthorn et al. [20] when investigating optimal control of epidemics in a scenario where no immunity to the disease is developed—is that the priority should be to protect the greater population of susceptible individuals, and that focusing on a subset of the population, rather than on the entire population, can make a significant difference [53]. We find that a larger allotment of vaccine can lead to the opposite result for a short period of time at the beginning of the time horizon, and we find that the high burden jurisdiction should be prioritized when it has more essential workers, as long as its infection level is increasing and remains higher than the jurisdiction with fewer essential workers. Our results also suggest that deviations from an *ad hoc* allocation rule based on population size are highly beneficial when one jurisdiction has a population with higher mortality rates.

We also show the value of complying to a travel restriction, as compliance leads to lower cumulative damages across both jurisdictions, regardless of whether immunity is permanent or temporary. The reduction in cumulative damages is particularly important for the jurisdiction with fewer infected individuals. Considering nonlinear damages due to an overload of health care systems [54; 55] and a corresponding varying death rate due to scarce intensive care unit beds [24], and other second-order problems such as consumption losses [56; 57; 58], excess mortality [59], and psychological distress [60] could further highlight the benefits of complying to travel restrictions.

Despite having to pay a workability cost for deviating from the *ad hoc* allocation, we show that it is still in the interest of the central planning agency (e.g. the central government) to deviate from this rule of thumb for a wide

range of values we considered; this result holds in all scenarios we considered in our analysis. We considered *ad hoc* allocation rules that favor “speed and workability” (put forward by NASEM [8], and WHO has a similar principle [9]). Other allocation rules are possible. For instance, in the base case of our paper, we assumed identical contact rates across jurisdictions. In turn, this implied that the movement within a given jurisdiction is assumed to be identical across jurisdictions. In practice, population mobility likely differs from one jurisdiction to another and an *ad hoc* allocation could be based on population mobility and contact structure. The methodology employed in this paper can investigate the trade-offs of other *ad hoc* rules and as a result, can offer potentially important information to policymakers that face the challenge of allocating scarce COVID-19 vaccines to their jurisdictions.

There are, however, important factors that have received significant attention in the literature that we should fully incorporate in future research. For example, the composition of the population of a jurisdiction is assumed to be homogeneous. To mimic the fact that the virus disproportionately affects elderly people [14] and/or people with pre-existing conditions [61], and also to mimic the fact that the risk of infection is highly occupation dependent [15], we simply assumed that the average case-fatality ratio and contact rate was higher in one jurisdiction. In practice, however, the composition of a population within a given jurisdiction is not homogeneous. Further research combining heterogeneity both across jurisdictions in the form of different disease burden and within jurisdictions in the form of different risk of complications and risk of infection could add additional valuable insights into the trade-offs inherent in these different allocations rules.

Finally, while our paper and most of the discussion revolves around the allocation of a vaccine, a similar allocation problem may arise if an antiviral drug were to become available (for a discussion on antiviral treatments for SARS-Cov-2, see [62]). Because drugs and vaccines have different goals—treating infected individuals and prophylaxis, respectively—the economic and public health trade-offs of different allocation rules may be unique to the type of pharmaceutical intervention. Future work considering the joint allocation question of antiviral drugs and vaccines could be valuable in understanding the trade-offs and complementarities between these different pharmaceutical interventions.

References

- [1] Emanuel E, Persad G, Upshur R, Thome B, Parker M, Glickman A, et al. Fair Allocation of Scarce Medical Resources in the Time of COVID-19. *The New England Journal of Medicine*. 2020.
- [2] Roope LS, Buckell J, Becker F, Candio P, Violato M, Sindelar JL, et al. How Should a Safe and Effective COVID-19 Vaccine be Allocated? Health Economists Need to be Ready to Take the Baton. *PharmacoEconomics-Open*. 2020:1–5.
- [3] Buckner JH, Chowell G, Springborn MR. Dynamic prioritization of COVID-19 vaccines when social distancing is limited for essential workers. *Proceedings of the National Academy of Sciences*. 2021;118(16).
- [4] Emanuel EJ, Persad G, Kern A, Buchanan A, Fabre C, Halliday D, et al. An ethical framework for global vaccine allocation. *Science (New York, NY)*. 2020:eabe2803.
- [5] Liu Y, Salwi S, Drolet B. Multivalue ethical framework for fair global allocation of a COVID-19 vaccine. *Journal of Medical Ethics*. 2020.
- [6] Yamey G, Schäferhoff M, Hatchett R, Pate M, Zhao F, McDade KK. Ensuring global access to COVID-19 vaccines. *The Lancet*. 2020;395(10234):1405–1406.
- [7] World Health Organization and others. WHO SAGE values framework for the allocation and prioritization of COVID-19 vaccination, 14 September 2020. World Health Organization; 2020.
- [8] National Academies of Sciences, Engineering, and Medicine. Discussion Draft of the Preliminary Framework for Equitable Allocation of COVID-19 Vaccine. 2020. <https://doi.org/10.17226/25914>.
- [9] Organization WH, et al. Fair allocation mechanism for COVID-19 vaccines through the COVAX Facility. Final working version-9 September. 2020.
- [10] Xia S, Zhang Y, Wang Y, Wang H, Yang Y, Gao GF, et al. Safety and immunogenicity of an inactivated SARS-CoV-2 vaccine, BBIBP-CorV: a randomised, double-blind, placebo-controlled, phase 1/2 trial. *The Lancet Infectious Diseases*. 2021;21(1):39–51.
- [11] Meo S, Bukhari I, Akram J, Meo A, Klonoff D. COVID-19 vaccines: comparison of biological, pharmacological characteristics and adverse effects of Pfizer/BioNTech and Moderna Vaccines. *Eur Rev Med Pharmacol Sci*. 2021;25(3):1663–1669.
- [12] Iacobucci G. Covid-19: Single vaccine dose is 33% effective against variant from India, data show. British Medical Journal Publishing Group; 2021.
- [13] Zaric GS, Brandeau ML. Optimal investment in a portfolio of HIV prevention programs. *Medical Decision Making*. 2001;21(5):391–408.

- [14] Verity R, Okell LC, Dorigatti I, Winskill P, Whittaker C, Imai N, et al. Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet infectious diseases*. 2020.
- [15] Baker MG, Peckham TK, Seixas NS. Estimating the burden of United States workers exposed to infection or disease: a key factor in containing risk of COVID-19 infection. *PLoS One*. 2020;15(4):e0232452.
- [16] Thomas LJ, Huang P, Yin F, Luo XI, Almquist ZW, Hipp JR, et al. Spatial heterogeneity can lead to substantial local variations in COVID-19 timing and severity. *Proceedings of the National Academy of Sciences*. 2020;117(39):24180–24187.
- [17] Polyakova M, Kocks G, Udalova V, Finkelstein A. Initial economic damage from the COVID-19 pandemic in the United States is more widespread across ages and geographies than initial mortality impacts. *Proceedings of the National Academy of Sciences*. 2020. Available from: <https://www.pnas.org/content/early/2020/10/19/2014279117>.
- [18] Bag R, Ghosh M, Biswas B, Chatterjee M. Understanding the spatio-temporal pattern of COVID-19 outbreak in India using GIS and India’s response in managing the pandemic. *Regional Science Policy & Practice*. 2020;12(6):1063–1103.
- [19] Brandeau ML, Zaric GS, Richter A. Resource allocation for control of infectious diseases in multiple independent populations: beyond cost-effectiveness analysis. *Journal of health economics*. 2003;22(4):575–598.
- [20] Rowthorn RE, Laxminarayan R, Gilligan CA. Optimal control of epidemics in metapopulations. *Journal of the Royal Society Interface*. 2009;6(41):1135–1144.
- [21] Zhou Y, Yang K, Zhou K, Liang Y. Optimal vaccination policies for an SIR model with limited resources. *Acta biotheoretica*. 2014;62(2):171–181.
- [22] Gersovitz M, Hammer JS. The economical control of infectious diseases. *The Economic Journal*. 2004;114(492):1–27.
- [23] Dangerfield CE, Vyska M, Gilligan CA. Resource Allocation for Epidemic Control Across Multiple Subpopulations. *Bulletin of mathematical biology*. 2019;81(6):1731–1759.
- [24] Acemoglu D, Chernozhukov V, Werning I, Whinston MD. Optimal targeted lockdowns in a multi-group SIR model. *NBER Working Paper*. 2020;27102.
- [25] Alvarez FE, Argente D, Lippi F. A simple planning problem for COVID-19 lockdown. *NBER Working Paper*. 2020;26981.
- [26] Levine-Tiefenbrun M, Yelin I, Katz R, Herzel E, Golan Z, Schreiber L, et al. Decreased SARS-CoV-2 viral load following vaccination. *medRxiv*. 2021.

- [27] Begon M, Bennett M, Bowers RG, French NP, Hazel S, Turner J. A clarification of transmission terms in host-microparasite models: numbers, densities and areas. *Epidemiology & Infection*. 2002;129(1):147–153.
- [28] Edridge AW, Kaczorowska J, Hoste AC, Bakker M, Klein M, Loens K, et al. Seasonal coronavirus protective immunity is short-lasting. *Nature medicine*. 2020;26(11):1691–1693.
- [29] Peng L, Yang W, Zhang D, Zhuge C, Hong L. Epidemic analysis of COVID-19 in China by dynamical modeling. *arXiv preprint arXiv:200206563*. 2020.
- [30] Yang Z, Zeng Z, Wang K, Wong SS, Liang W, Zanin M, et al. Modified SEIR and AI prediction of the epidemics trend of COVID-19 in China under public health interventions. *Journal of Thoracic Disease*. 2020;12(3):165.
- [31] Hou C, Chen J, Zhou Y, Hua L, Yuan J, He S, et al. The effectiveness of quarantine of Wuhan city against the Corona Virus Disease 2019 (COVID-19): A well-mixed SEIR model analysis. *Journal of medical virology*. 2020.
- [32] Pandey G, Chaudhary P, Gupta R, Pal S. SEIR and Regression Model based COVID-19 outbreak predictions in India. *arXiv preprint arXiv:200400958*. 2020.
- [33] Prem K, Liu Y, Russell TW, Kucharski AJ, Eggo RM, Davies N, et al. The effect of control strategies to reduce social mixing on outcomes of the COVID-19 epidemic in Wuhan, China: a modelling study. *The Lancet Public Health*. 2020.
- [34] Radulescu A, Cavanagh K. Management strategies in a SEIR model of COVID 19 community spread. *arXiv preprint arXiv:200311150*. 2020.
- [35] Roda WC, Varughese MB, Han D, Li MY. Why is it difficult to accurately predict the COVID-19 epidemic? *Infectious Disease Modelling*. 2020.
- [36] Bjørnstad ON, Shea K, Krzywinski M, Altman N. Modeling infectious epidemics. *Nature methods*. 2020;17:455–456.
- [37] Bjørnstad O, Shea K, Krzywinski M, Altman N. The SEIRS model for infectious disease dynamics. *Nature Methods*. 2020;17(6):557–558.
- [38] Stutt RO, Retkute R, Bradley M, Gilligan CA, Colvin J. A modelling framework to assess the likely effectiveness of facemasks in combination with ‘lock-down’ in managing the COVID-19 pandemic. *Proceedings of the Royal Society A*. 2020;476(2238):20200376.
- [39] Bertozzi AL, Franco E, Mohler G, Short MB, Sledge D. The challenges of modeling and forecasting the spread of COVID-19. *Proceedings of the National Academy of Sciences*. 2020;117(29):16732–16738.

- [40] Burton J, Billings L, Cummings DA, Schwartz IB. Disease persistence in epidemiological models: the interplay between vaccination and migration. *Mathematical biosciences*. 2012;239(1):91–96.
- [41] Chen M, Li M, Hao Y, Liu Z, Hu L, Wang L. The introduction of population migration to SEIAR for COVID-19 epidemic modeling with an efficient intervention strategy. *Information Fusion*. 2020;64:252–258.
- [42] Goldman SM, Lightwood J. Cost optimization in the SIS model of infectious disease with treatment. *The BE Journal of Economic Analysis & Policy*. 2002;2(1).
- [43] Barrett S, Hoel M. Optimal disease eradication. *Environment and Development Economics*. 2007:627–652.
- [44] Sivaraman NK, Gaur M, Bajjal S, Rupesh CV, Muthiah SB, Sheth A. Exo-SIR: An Epidemiological Model to Analyze the Impact of Exogenous Infection of COVID-19 in India. *arXiv preprint arXiv:200806335*. 2020.
- [45] World Tourism Organization. International Tourism Down 70% As Travel Restrictions Impact All Regions. Madrid, Spain: United Nations World Tourism Organization, 2020. <https://www.unwto.org/news/international-tourism-down-70-as-travel-restrictions-impact-all-regions> (accessed October 27, 2020); 2020.
- [46] Environmental Protection Agency. What value of statistical life does EPA use? Washington, DC: Environmental Protection Agency, 2020. <https://www.epa.gov/environmental-economics/mortality-risk-valuation#whatvalue> (accessed October 27, 2020); 2020.
- [47] Nurchis MC, Pascucci D, Sapienza M, Villani L, D’Ambrosio F, Castrini F, et al. Impact of the Burden of COVID-19 in Italy: Results of Disability-Adjusted Life Years (DALYs) and Productivity Loss. *International Journal of Environmental Research and Public Health*. 2020;17(12):4233.
- [48] Ryan D, Toews C, Sanchirico JN, Armsworth PR. Implications of policy adjustment costs for fisheries management. *Natural Resource Modeling*. 2017;30(1):74–90.
- [49] Kling DM, Sanchirico JN, Wilen JE. Bioeconomics of managed relocation. *Journal of the Association of Environmental and Resource Economists*. 2016;3(4):1023–1059.
- [50] Castonguay FM, Sokolow SH, De Leo GA, Sanchirico JN. Cost-effectiveness of combining drug and environmental treatments for environmentally transmitted diseases. *Proceedings of the Royal Society B*. 2020;287(1933):20200966.
- [51] Ohmit SE, Thompson MG, Petrie JG, Thaker SN, Jackson ML, Belongia EA, et al. Influenza vaccine effectiveness in the 2011–2012 season: protection against each circulating virus and the effect of prior vaccination on estimates. *Clinical infectious diseases*. 2014;58(3):319–327.

- [52] England PH. SARS-CoV-2 variants of concern and variants under investigation in England. technical briefing 17. 2021.
- [53] Duijzer LE, van Jaarsveld WL, Wallinga J, Dekker R. Dose-optimal vaccine allocation over multiple populations. *Production and Operations Management*. 2018;27(1):143–159.
- [54] Phadke I, McKee A, Conway JM, Shea K. Analysing how changes in the health status of healthcare workers affects epidemic outcomes. *Epidemiology and Infection*. 2021;149:e42.
- [55] Verelst F, Kuylen E, Beutels P. Indications for healthcare surge capacity in European countries facing an exponential increase in coronavirus disease (COVID-19) cases, March 2020. *Eurosurveillance*. 2020;25(13):2000323.
- [56] Baker SR, Farrokhnia RA, Meyer S, Pagel M, Yannelis C. How does household spending respond to an epidemic? consumption during the 2020 COVID-19 pandemic. *National Bureau of Economic Research*; 2020.
- [57] Andersen AL, Hansen ET, Johannesen N, Sheridan A. Consumer responses to the COVID-19 crisis: Evidence from bank account transaction data. Available at SSRN 3609814. 2020.
- [58] Coibion O, Gorodnichenko Y, Weber M. The cost of the COVID-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. *National Bureau of Economic Research*; 2020.
- [59] Vestergaard LS, Nielsen J, Richter L, Schmid D, Bustos N, Braeye T, et al. Excess all-cause mortality during the COVID-19 pandemic in Europe—preliminary pooled estimates from the EuroMOMO network, March to April 2020. *Eurosurveillance*. 2020;25(26):2001214.
- [60] Pfefferbaum B, North CS. Mental health and the Covid-19 pandemic. *New England Journal of Medicine*. 2020.
- [61] Ssentongo P, Ssentongo AE, Heilbrunn ES, Ba DM, Chinchilli VM. Association of cardiovascular disease and 10 other pre-existing comorbidities with COVID-19 mortality: A systematic review and meta-analysis. *PloS one*. 2020;15(8):e0238215.
- [62] Hu B, Guo H, Zhou P, Shi ZL. Characteristics of SARS-CoV-2 and COVID-19. *Nature Reviews Microbiology*. 2020:1–14.

Appendices

Contents

A	Parameterization	2
A.1	Epidemiological Model	2
A.2	Economic Model	2
A.3	Parameter Levels	3
B	Optimization	4
B.1	Boundary Conditions	4
B.2	Nonnegativity and Upper-Bound Constraints	4
B.3	Capacity Constraints of the Pharmaceutical Interventions	4
B.4	Numerical Methods	5
C	Figures	6
C.1	Base Case: Homogeneous Demographic Characteristics	6
C.1.1	Compliance and Noncompliance to the Travel Restrictions	6
C.1.2	Vaccine Capacity Constraints when Immunity is Permanent	7
C.1.3	Vaccine Capacity Constraints when Immunity is Temporary	8
C.1.4	Permanent vs Temporary Immunity	9
C.1.5	Cumulative Infection Levels	10
C.2	Heterogeneous Demographic Characteristics	12
C.2.1	Heterogeneous Case-Fatality Ratio	12
C.2.2	Heterogeneous Contact Rate	13
C.3	Robustness of Spatial allocations	14
C.3.1	Base Case: Homogeneous Demographic Characteristics	14
C.3.2	Heterogeneous Case-Fatality Ratio	15
C.3.3	Heterogeneous Contact Rate	16
C.4	Sensitivity Analyses	17
C.4.1	Workability Cost	17
C.4.2	Vaccine Effectiveness	19

A Parameterization

A.1 Epidemiological Model

According to Diekmann et al. [1], the basic reproduction ratio R_0 of any disease is given by the expected number of secondary infection caused by a typical infected individual over its entire infectious period, at a disease-free equilibrium. In the most basic epidemiological model, the R_0 is simply given by the contact rate multiplied by the mean infectious period. When considering more complex models—as the two-jurisdiction SEIR model in this paper—one needs to use the next-generation matrix and find its dominant eigenvalue to find the R_0 [1]. Denote two matrices by F and V , and let the ij^{th} element in F represents the rate at which infected individuals in population j produce new infections in population i , and the ij^{th} element in V represents the transition rate between $(i \neq j)$, or out of $(i = j)$, infectious compartments [2]; the next-generation matrix is equal to $-FV^{-1}$. In the model presented in this paper,

$$F = \begin{pmatrix} 0 & \beta_{11} & 0 & \beta_{12} \\ 0 & 0 & 0 & 0 \\ 0 & \beta_{21} & 0 & \beta_{22} \\ 0 & 0 & 0 & 0 \end{pmatrix} \text{ and } V = \begin{pmatrix} -\sigma & 0 & 0 & 0 \\ \sigma & -(\gamma + \varphi_1) & 0 & 0 \\ 0 & 0 & -\sigma & 0 \\ 0 & 0 & \sigma & -(\gamma + \varphi_2) \end{pmatrix}$$

where the four rows of F and V refer to the E_1 , I_1 , E_2 and I_2 equations, respectively. Note that both matrices F and V are derived under the assumption of introducing a single exposed individuals in an otherwise susceptible population (for more details on how to construct the next-generation matrix in a SEIR model, see [3]). Given we assume that $\beta_{11} = \beta_{22} = \beta_{ii}$ and $\beta_{12} = \beta_{21} = \beta_{ij}$, and when in our main analysis we let $\varphi_1 = \varphi_2 = \varphi$, the basic reproduction ratio of our model simplifies to,

$$R_0 = \frac{\beta_{ii} + \beta_{ij}}{\gamma + \varphi}$$

for $i = 1, 2$, $j = 1, 2$, and $i \neq j$. We set the basic reproduction ratio $R_0 = 1.43$, according to estimates of the R_0 from Li et al. [4] and using estimates of the effect of nonpharmaceutical interventions on the R_0 from Tian et al. [5]. We assume a mean recovery period ($\frac{1}{\gamma}$) of 5 days [6], and a case-fatality ratio of 1.78% (adjusted for misreporting, see [7]) to calibrate the rate of disease induced mortality, φ . Parameters β_{ii} and β_{ij} are then calibrated assuming what Tian et al. [5] call a “medium effect of the [nonpharmaceutical] control” when there is compliance to travel restrictions, and a “lower effect of the [nonpharmaceutical] control” when there is no compliance to travel restrictions (for evidence of structural changes in mobility following the COVID-19 lockdown, see [8]); this yields $R_0 \approx 1.4$ with compliance to travel restrictions, and $R_0 \approx 2.1$ when there is no compliance to travel restrictions. The mean latency period ($\frac{1}{\sigma}$), which one needs to know to calculate matrix V even though it does not appear in the basic reproduction ratio, is assumed to last 3 days [6].

A.2 Economic Model

To quantify damages, we use the value of statistical life recommended by the Environmental Protection Agency.¹ The disability weight² associated with COVID-19 infection is assumed to be equivalent to a lower respiratory tract infection, which is a disability weight of $w = 0.133$ on a scale from zero (perfect health) to one (death).³ This disability weight thus allows for a comparison between the individuals that are infected with the disease but do not die, and the individuals that die from its complications.

Expenditures related to the pharmaceutical intervention are based off estimates of vaccine costs. Numerous governments around the world, including the U.S. federal government, have contracted biotech companies producing COVID-19 vaccines; governments pay money in exchange of a guaranteed number of doses of COVID-19 vaccines. These estimates and the prices of current influenza vaccine turn out to be approximately 20 U.S. dollars per dose, with two doses per individual; this is the value we chose in our analysis.⁴

¹See “What value of statistical life does EPA use?” from the U.S. Environmental Protection Agency [9].

²According to the World Health Organization: “A disability weight is a weight factor that reflects the severity of the disease on a scale from 0 (perfect health) to 1 (equivalent to death).” See: https://www.who.int/healthinfo/global_burden_disease/daly_disability_weight/en/.

³For more details on how COVID-19’s disability resembles lower respiratory tract infections, see [10].

⁴For COVID-19 vaccine prices, see: <https://www.npr.org/sections/health-shots/2020/08/06/899869278/prices-for-covid-19-vaccines-are-starting-to-come-into-focus>. For a comparison with influenza vaccine prices, see <https://www.cdc.gov/vaccines/programs/vfc/awardees/vaccine-management/price-list/index.html>.

The value of the workability cost⁵ is based on a certain proportion of the value of statistical life; in the base case, we assume it to be 3 orders of magnitude smaller. All costs in the model are assumed to be discounted at a 1.5% annual rate (see [12] for a discussion about discounting health-related expenditures).

A.3 Parameter Levels

Table A1 below summarizes the main set of parameter values we used in the numerical simulation.

Parameters	Level	Definition
β_{ii}	8.86	Transmission rate within a given state (month ⁻¹). ⁶
β_{ij}	4.36	Transmission rate across states (month ⁻¹). ⁶
σ	10.14	Rate at which infected individuals become infectious (month ⁻¹). ⁷
γ	6.08	Rate of recovery (month ⁻¹). ⁷
ω	0.17	Rate at which immunity is lost (month ⁻¹). ⁸
φ	0.11	Rate of disease induced mortality (month ⁻¹). ⁹
w	0.13	Disability weight associated with the disease (unitless). ¹⁰
q_V	0.65	Efficiency of vaccines (proportion). ¹¹
r	0.0013	Discount rate (month ⁻¹). ¹²
c_V	40	Cost of treating one individual via vaccine (US Dollars). ¹³
c_A	10×10^3	Workability cost (US Dollars). ¹⁴
c	10×10^6	Value of statistical life (US Dollars). ¹⁵

Table A1: Parameter levels used in the numerical simulation.

⁵Inspired by the paper of Ryan et al. [11] where the authors show the implications of policy adjustment costs for fisheries management

⁶Calibrated using a R_0 estimate from Li et al. [4] and estimates of effects of nonpharmaceutical interventions from Tian et al. [5]; this yields a R_0 of approximately 1.4 when there is compliance to travel restrictions and to match a R_0 of approximately 2.1 when there is no compliance to travel restrictions; these two values representing respectively a “medium” and “low” effect of the nonpharmaceutical intervention.

⁷Using estimates from Davies et al. [6]; this represents a 3-day latency period and a 5-day recovery period.

⁸Representing a 6-month immunity period in the scenarios where we assume immunity is not permanent; based on [13].

⁹Calibrated by using a case-fatality rate of 1.78%. Adjusted for mis- and under-reporting; see [7].

¹⁰Representing the disability associated with severe lower respiratory tract infections because, to our knowledge, there are no official disability estimates associated with COVID-19; see [10].

¹¹Following Buckner et al. [14], we base this parameter value on the efficiency of the influenza vaccine [15]. Note that the Department of Health and Human Services, Food and Drug Administration (FDA), and Center for Biologics Evaluation and Research (CBER) [16] requires that a future COVID-19 vaccine must have an effectiveness of at least 50%.

¹²Based on results from John et al. [12] that suggest a yearly discount rate between 0.3% and 1.5% for health related expenditures; we chose a 1.5% annual discount rate in the main set of results. This gives a monthly discount rate of $r = 0.0013$.

¹³Assuming an individual requires two doses; based on current agreements between the U.S. federal government and biotech companies; see <https://www.npr.org/sections/health-shots/2020/08/06/899869278/prices-for-covid-19-vaccines-are-starting-to-come-into-focus>. For a list of current vaccine prices, and particularly the price of the influenza vaccine, see <https://www.cdc.gov/vaccines/programs/vfc/awardees/vaccine-management/price-list/index.html>.

¹⁴Value based on a certain proportion of the value of statistical life, c ; in the base case we assume it is 2 orders of magnitude smaller.

¹⁵Represents a value of statistical life of 10M U.S. dollars. Based on the value of a statistical life that the U.S. Environmental Protection Agency [9] uses: approximately \$7.4 million (\$2006) which is equivalent to approximately \$9.54 million (\$2020).

B Optimization

B.1 Boundary Conditions

To yield the initial conditions of the optimal control problem, we calibrated the model using the above parameter values and simulated out a COVID-19 outbreak in two identical jurisdictions, where we assumed there was one exposed individual in an otherwise entirely susceptible population of 10 million individuals. We assumed that both jurisdictions undertook nonpharmaceutical interventions that had a “medium effect” on the basic reproduction ratio [5] (i.e. that there was perfect compliance to travel restrictions). After simulating out the disease dynamics for a period of eight months and two weeks, and eight months and three weeks for Jurisdiction 1 and Jurisdiction 2 respectively, the initial conditions yield were:

Jurisdiction	N_i	S_i	E_i	I_i	R_i
Jurisdiction 1	1	0.9074	0.0103	0.0143	0.0667
Jurisdiction 2	1	0.8662	0.0138	0.0196	0.0986

Table A2: Initial conditions of the numerical simulation.

We assume that the terminal conditions (i.e. the conditions on state variables in $t = T$, the final time period) are free to be optimally determined. Formally, the initial and terminal conditions of the ten state variables are such that:

$$S_i(0), E_i(0), I_i(0), R_i(0), \text{ and } N_i(0) \text{ are given for } i = 1, 2; \quad (\text{A1a})$$

$$S_i(T), E_i(T), I_i(T), R_i(T), \text{ and } N_i(T) \text{ are free for } i = 1, 2. \quad (\text{A1b})$$

B.2 Nonnegativity and Upper-Bound Constraints

State variables S_i , E_i , I_i , R_i , and N_i for $i = 1, 2$ are subject to nonnegativity and physical constraints. Formally:

$$0 \leq S_i \leq N_i \leq 1 \text{ for } i = 1, 2; \quad (\text{A2a})$$

$$0 \leq E_i \leq N_i \leq 1 \text{ for } i = 1, 2; \quad (\text{A2b})$$

$$0 \leq I_i \leq N_i \leq 1 \text{ for } i = 1, 2; \quad (\text{A2c})$$

$$0 \leq R_i \leq N_i \leq 1 \text{ for } i = 1, 2; \quad (\text{A2d})$$

$$S_i + E_i + I_i + R_i = N_i \leq 1 \text{ for } i = 1, 2. \quad (\text{A2e})$$

Control variables are modelled as direct controls (see examples in [17; 18; 19]) and can be interpreted as a reduction in the number of susceptible individual in a given time period (i.e. a month). Formally, the constraints on the control variables are given by:

$$0 \leq u_{V_i} \leq S_i \text{ for } i = 1, 2. \quad (\text{A3})$$

Because of a limited supply of vaccines (see details below), the physical upper-bound on constraints (A3) will only be binding when capacity constraint is nonbinding. When this occurs, it means that there are fewer susceptible individuals than there are available vaccines.

B.3 Capacity Constraints of the Pharmaceutical Interventions

For completeness, we also include the capacity constraints already mentioned in the main paper. In addition to the physical constraints on the control variables, the aim of our paper is to study how to allocate a limited allotment of vaccines. Hence, the control variables are also subject to

$$u_{V_1} + u_{V_2} \leq \bar{u}_V; \quad (\text{A4})$$

when the central planning agency decides to potentially deviate from the *ad hoc* allocation of vaccine. Conversely, the *ad hoc* constraints are:

$$u_{V_i} \leq \left(\frac{N_i}{N_1 + N_2} \right) \bar{u}_V \quad \text{for } i = 1, 2. \quad (\text{A5})$$

As mentioned in the main paper, the total available quantity of vaccine ($\bar{u}_V$) represents a certain percentage (5%, 10%, or 15%) of the total population size.

B.4 Numerical Methods

Pseudospectral collocation approximates the continuous time optimal control model with a constrained nonlinear programming problem (see [20; 21; 22; 23] for other applications of this technique). The dynamic controls to our problem—i.e. the vaccine allocation—are approximated by a polynomial of degree n (determined by the number of collocation points) over a period from $t = 0$ (date at which the vaccine starts to be administered) to $t = T$ (assumed to be four months after the vaccine administration) [24]. The residual error of the constraints is minimized by the algorithm at the n collocation points, where n is chosen to have a reasonable speed of convergence to a solution and a low numerical error. Here, we chose 60 collocations points. In this sort of problem, the main advantage of this approach over more usual methods to solve such two-point boundary problems, such as shooting methods, is that nonnegativity constraints (e.g. on the number of infected individuals) and upper-bound constraints (mimicking e.g. vaccine capacity constraints) on state and control variables can be directly incorporated in the problem [25]. This method thus allows us to find optimal solutions that may lay on the boundary of the control set for a certain period of time. For COVID-19 vaccines, this is likely due to the scarcity of the supply of vaccine in the short-term. Another advantage of this method is the ability to deal with large-scale dynamical systems, such as the one presented here with ten state variables and two control variables. The solution was found using TOMLAB (v. 8.4) [26; 27] and the accompanying PROPT toolbox [28]. The approximate nonlinear programming problem is solved using general-purpose nonlinear optimization packages (e.g. KNITRO, SNOPT and NPSOL).

C Figures

C.1 Base Case: Homogeneous Demographic Characteristics

C.1.1 Compliance and Noncompliance to the Travel Restrictions

Figure A1: **Permanent immunity with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether there is compliance to travel restrictions (panels A and C) or not (panels B and D) for the case where the vaccine capacity constraint is 5% and immunity lasts six months.

C.1.2 Vaccine Capacity Constraints when Immunity is Permanent

Figure A2: **Permanent immunity and compliance to travel restrictions with 5%, 10%, and 15% vaccine capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether capacity is 5% (panels A and D), 10% (panels B and E), or 15% (panels C and F), for the case where immunity is permanent and there is compliance to travel restrictions.

C.1.3 Vaccine Capacity Constraints when Immunity is Temporary

Figure A3: **Temporary immunity and compliance to travel restrictions with 5%, 10%, and 15% vaccine capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether capacity is 5% (panels A and D), 10% (panels B and E), or 15% (panels C and F), for the case where immunity lasts six months and there is compliance to travel restrictions.

Figure A4: **Temporary immunity and noncompliance to travel restrictions with 5%, 10%, and 15% vaccine capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether capacity is 5% (panels A and D), 10% (panels B and E), 15% (panels C and F), for the case where immunity lasts six months and there is no compliance to travel restrictions.

C.1.4 Permanent vs Temporary Immunity

Figure A5: **Compliance to travel restrictions with permanent and temporary immunity.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether immunity is permanent (panels A and C) or lasts six months (panels B and D) for the case where the vaccine capacity constraint is 10% and there is compliance to travel restrictions.

Figure A6: **Noncompliance to travel restrictions with permanent and temporary immunity.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether immunity is permanent (panels A and C) or lasts six months (panels B and D) for the case where the vaccine capacity constraint is 10% and there is no compliance to travel restrictions.

C.1.5 Cumulative Infection Levels

Figure A7: **Epidemiological outcomes under different scenarios with a low vaccine supply.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-vaccine case for different immunity-travel restrictions scenarios and for when vaccine capacity is 5%.

Figure A8: **Epidemiological outcomes under different scenarios with a high vaccine supply.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-vaccine case for different immunity-travel restrictions scenarios and for when vaccine capacity is 15%.

C.2 Heterogeneous Demographic Characteristics

C.2.1 Heterogeneous Case-Fatality Ratio

Figure A9: **Vaccine allocation with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether there is compliance to travel restrictions (panels A and C) or not (panels B and D) for the case where the vaccine capacity constraint is 10%, immunity is permanent, and where the heterogeneity in the system comes from a varying case-fatality ratio (State 1 has a case-fatality ratio 1% higher than State 2).

Figure A10: **Vaccine allocation with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether there is compliance to travel restrictions (panels A and C) or not (panels B and D) for the case where the vaccine capacity constraint is 10%, immunity lasts six months, and where the heterogeneity in the system comes from a varying case-fatality ratio (State 1 has a case-fatality ratio 1% higher than State 2).

C.2.2 Heterogeneous Contact Rate

Figure A11: **Vaccine allocation with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether there is compliance to travel restrictions (panels A and C) or not (panels B and D) for the case where the vaccine capacity constraint is 10%, immunity is permanent, and where the heterogeneity in the system comes from a varying contact rate (State 2 has a higher contact rate than State 1).

Figure A12: **Vaccine allocation with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether there is compliance to travel restrictions (panels A and C) or not (panels B and D) for the case where the vaccine capacity constraint is 10%, immunity lasts six months, and where the heterogeneity in the system comes from a varying contact rate (State 2 has a higher contact rate than State 1).

C.3 Robustness of Spatial allocations

C.3.1 Base Case: Homogeneous Demographic Characteristics

Figure A13: **Robustness of epidemiological and economic outcomes under different scenarios.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 5%.

Figure A14: **Robustness of epidemiological and economic outcomes under different scenarios.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 15%.

C.3.2 Heterogeneous Case-Fatality Ratio

Figure A15: **Robustness of epidemiological and economic outcomes under different scenarios when the source of heterogeneity is the case-fatality ratio.** Percentage change in expenditures (y -axis) and percentage change in cumulative cases (x -axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 10%.

Figure A16: **Robustness of epidemiological and economic outcomes under different scenarios when the source of heterogeneity is the case-fatality ratio.** Percentage change in expenditures (y -axis) and percentage change in cumulative damages (x -axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 10%. Note that compared to Figure A15, the use of cumulative damages in this figure gives a more accurate depiction of the situation because cases across jurisdictions are not homogeneous when the case-fatality ratio is different.

C.3.3 Heterogeneous Contact Rate

Figure A17: **Robustness of epidemiological and economic outcomes under different scenarios when the source of heterogeneity is the contact rate.** Percentage change in expenditures (y -axis) and percentage change in cumulative cases (x -axis) from the optimal allocation for different immunity-travel restrictions scenarios and for when vaccine capacity is 10%.

C.4 Sensitivity Analyses

C.4.1 Workability Cost

As mentioned above, imposing the *ad hoc* rule *ex ante* implicitly means that the central planning agency is essentially assuming that the cost of deviating from the *ad hoc* allocation is infinite. In practice, the workability cost is hard to quantify, because it depends on logistical, political, and cultural factors. It does however seem reasonable to assume, as we did in the paper, that the cost is finite. We investigate the sensitivity of our results by solving for the optimal vaccine allocation over time with levels lower and higher than the base case parameter in the paper. We summarize these results by plotting the variance¹⁶ of the optimal deviation in each time period from the *ad hoc* vaccine allocation (in blue; Figure A18 panels A, B, C, and D), and the difference in cumulative cases between the optimal and *ad hoc* allocation (in red; Figure A18 panels A, B, C, and D) as we vary the scale of the workability cost. Mathematically, as the workability cost approaches zero, the optimal control problem becomes linear in the controls, which implies that there is no adjustment cost associated with changing the allocation. Often times this can lead to extreme solutions (allocation goes to one state for a time period and then the other state, and so on).

Given the behavior and nature of the problem, therefore, we expect that at lower values of the workability cost parameter we will find higher variance of the deviation. This, in turn results in a higher performance of the optimal allocation relative to the *ad hoc* in terms of reduction in cumulative cases. When we increase the workability cost parameter, the cost parameter will eventually be on the same magnitude as the VSL (Figure A18 black line represents the VSL). When we reach levels this high, the optimal allocation converges towards the *ad hoc* and any differences in cumulative cases disappear.

We also show how amount of funds allocated to the workability cost over time compare to expenditures on the total vaccine cost (Figure A18 panels E, F, G, and H). If this ratio exceeds one, the planner is spending on aggregate more to deviate from the *ad hoc* than on treatments. These panels show that at low levels of the workability cost parameter, the total workability costs are small relative to the total vaccine costs. As the workability cost parameter increases, however, the total workability costs become more and more important relative to the total vaccine cost. Eventually, these costs begin to dominate the planners objective and the deviation between the *ad hoc* and optimal goes to zero.

¹⁶The variance is calculated as $\text{Var}\left(\frac{\text{Optimal Vaccine} - \text{Ad Hoc Vaccine}}{\text{Ad Hoc Vaccine}}\right)$. Note that the variance of the optimal deviation from the *ad hoc* is identical in absolute and relative terms across jurisdictions.

Figure A18: **Sensitivity of optimal allocations and epidemiological and economic outcomes when varying the workability cost parameter.** The variance of the optimal deviation in percentage (in blue; panels A, B, C, and D) represents an aggregate measure of the optimal deviation from the *ad hoc* allocation. The difference in cumulative cases between the optimal and *ad hoc* allocations (in red; panels A, B, C, and D) represents in percentage terms how well the optimal allocation outperforms the *ad hoc* allocation. The total workability cost over the total vaccine cost (panels E, F, G, and H) represents how many times more the total workability costs are relative to the total vaccine costs. The dotted vertical line represents the base case value of the workability cost parameter ($1e^4$), while the full vertical line represents the value of statistical life ($1e^7$).

C.4.2 Vaccine Effectiveness

The base case parameter for vaccine effectiveness we utilized in the paper is based on estimates of the influenza vaccine [15]; see Appendix A for more details. Recent evidence from the COVID-19 vaccines suggest that effectiveness could be considerably higher. As a result, we investigate how a more effective vaccine would affect the nature of our results. We find that the more effective a vaccine is, the more a central planner would want to deviate from the *ad hoc* allocation (in blue; Figure A19 panels A, B, C, and D). As a result of this greater deviation, we see a larger difference in terms of the reduction in cumulative cases (in red; Figure A19 panels A, B, C, and D). Because a higher effectiveness results in a greater deviation, then, everything else equal, the total workability costs are increased relative to the total vaccine costs (Figure A19 panels E, F, G, and H). The differences are more stark in a world where there is compliance to travel restrictions, as noncompliance blurs the spatial heterogeneity across the jurisdictions leading in general to allocations similar to the *ad hoc*.

Figure A19: **Sensitivity of optimal allocations, and of epidemiological and economic outcomes when varying the effectiveness of the vaccine.** The variance of the optimal deviation in percentage (in blue; panels A, B, C, and D) represents an aggregate measure of the optimal deviation from the *ad hoc* allocation. The difference in cumulative cases between the optimal and *ad hoc* allocations (in red; panels A, B, C, and D) represents in percentage terms how well the optimal allocation outperforms the *ad hoc* allocation. The total workability cost over the total vaccine cost (panels E, F, G, and H) represents how many times more the total workability costs are relative to the total vaccine costs. The dotted vertical line in the plots represents the base case value of the vaccine effectiveness (0.65).

References

- [1] Diekmann O, Heesterbeek JAP, Metz JA. On the definition and the computation of the basic reproduction ratio R_0 in models for infectious diseases in heterogeneous populations. *Journal of Mathematical Biology*. 1990;28(4):365–382.
- [2] Garchitorena A, Sokolow S, Roche B, Ngonghala C, Jocque M, Lund A, et al. Disease ecology, health and the environment: a framework to account for ecological and socio-economic drivers in the control of neglected tropical diseases. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2017;372(1722):20160128.
- [3] Diekmann O, Heesterbeek J, Roberts MG. The construction of next-generation matrices for compartmental epidemic models. *Journal of the Royal Society Interface*. 2010;7(47):873–885.
- [4] Li Q, Guan X, Wu P, Wang X, Zhou L, Tong Y, et al. Early transmission dynamics in Wuhan, China, of novel coronavirus-infected pneumonia. *New England Journal of Medicine*. 2020.
- [5] Tian H, Liu Y, Li Y, Wu CH, Chen B, Kraemer MU, et al. An investigation of transmission control measures during the first 50 days of the COVID-19 epidemic in China. *Science*. 2020;368(6491):638–642.
- [6] Davies NG, Klepac P, Liu Y, et al. Age-dependent effects in the transmission and control of COVID-19 epidemics. *Nature Medicine*. 2020;26(8):1205–1211.
- [7] Abdollahi E, Champredon D, Langley JM, Galvani AP, Moghadas SM. Temporal estimates of case-fatality rate for COVID-19 outbreaks in Canada and the United States. *CMAJ*. 2020.
- [8] Schlosser F, Maier BF, Jack O, Hinrichs D, Zachariae A, Brockmann D. COVID-19 lockdown induces disease-mitigating structural changes in mobility networks. *Proceedings of the National Academy of Sciences*. 2020.
- [9] Environmental Protection Agency. What value of statistical life does EPA use? Washington, DC: Environmental Protection Agency, 2020. <https://www.epa.gov/environmental-economics/mortality-risk-valuation#whatvalue> (accessed October 27, 2020); 2020.
- [10] Nurchis MC, Pascucci D, Sapienza M, Villani L, D’Ambrosio F, Castrini F, et al. Impact of the Burden of COVID-19 in Italy: Results of Disability-Adjusted Life Years (DALYs) and Productivity Loss. *International Journal of Environmental Research and Public Health*. 2020;17(12):4233.
- [11] Ryan D, Toews C, Sanchirico JN, Armsworth PR. Implications of policy adjustment costs for fisheries management. *Natural Resource Modeling*. 2017;30(1):74–90.
- [12] John J, Koerber F, Schad M. Differential discounting in the economic evaluation of healthcare programs. *Cost Effectiveness and Resource Allocation*. 2019;17(1):29.
- [13] Edridge AW, Kaczorowska J, Hoste AC, Bakker M, Klein M, Loens K, et al. Seasonal coronavirus protective immunity is short-lasting. *Nature medicine*. 2020;26(11):1691–1693.
- [14] Buckner JH, Chowell G, Springborn MR. Dynamic prioritization of COVID-19 vaccines when social distancing is limited for essential workers. *Proceedings of the National Academy of Sciences*. 2021;118(16).
- [15] Ohmit SE, Thompson MG, Petrie JG, Thaker SN, Jackson ML, Belongia EA, et al. Influenza vaccine effectiveness in the 2011–2012 season: protection against each circulating virus and the effect of prior vaccination on estimates. *Clinical infectious diseases*. 2014;58(3):319–327.
- [16] U S Department of Health and Human Services, Food and Drug Administration (FDA), and Center for Biologics Evaluation and Research (CBER). Development and licensure of vaccines to prevent COVID-19: guidance for industry. June, 2020. Food and Drug Administration, 5630 Fishers Lane, Rm. 1061, Rockville, MD 20852; 2020.
- [17] Rowthorn RE, Laxminarayan R, Gilligan CA. Optimal control of epidemics in metapopulations. *Journal of the Royal Society Interface*. 2009;6(41):1135–1144.
- [18] Goldman SM, Lightwood J. Cost optimization in the SIS model of infectious disease with treatment. *The BE Journal of Economic Analysis & Policy*. 2002;2(1).
- [19] Barrett S, Hoel M. Optimal disease eradication. *Environment and Development Economics*. 2007:627–652.

- [20] Sanchirico JN, Springborn M. How to get there from here: ecological and economic dynamics of ecosystem service provision. *Environmental and Resource Economics*. 2011;48(2):243–267.
- [21] Kling DM, Sanchirico JN, Wilen JE. Bioeconomics of managed relocation. *Journal of the Association of Environmental and Resource Economists*. 2016;3(4):1023–1059.
- [22] Castonguay FM, Sokolow SH, De Leo GA, Sanchirico JN. Cost-effectiveness of combining drug and environmental treatments for environmentally transmitted diseases. *Proceedings of the Royal Society B*. 2020;287(1933):20200966.
- [23] Castonguay FM, Lasserre P. L’exploitation de ressources naturelles non renouvelables en asymétrie d’information. *L’Actualité économique*. 2019;95(2-3).
- [24] Garg D, Patterson M, Hager WW, Rao AV, Benson DA, Huntington GT. A unified framework for the numerical solution of optimal control problems using pseudospectral methods. *Automatica*. 2010;46(11):1843–1851.
- [25] Judd KL. *Numerical methods in economics*. MIT press; 1998. Cambridge, Massachusetts.
- [26] Holmström K. Practical optimization with the tomlab environment in matlab. In: *Proceedings of the 42nd SIMS Conference*, Sep 15, 2001. Citeseer; 2001. p. 89–108. Porsgrunn, Norway: Telemark University College.
- [27] Holmström K, Göran AO, Edvall MM. *Users guide for TOMLAB/SNOPT*. Mälardalen University, Department of Mathematics and Physics, Västerås, Sweden. 2008.
- [28] Rutquist PE, Edvall MM. Propt-Matlab optimal control software. Tomlab Optimization Inc. 2010;260(1).