Spatial Allocation of Scarce Vaccine and Antivirals for COVID-19*†‡

François M. Castonguay,§ Julie C. Blackwood,¶ Emily Howerton,‖
Katriona Shea,‖ Charles Sims,** and James N. Sanchirico††

December 19, 2020

Abstract

Although the COVID-19 disease burden is heterogeneous across space, the U.S. National Academies of Sciences, Engineering, and Medicine recommends an equitable spatial allocation of pharmaceutical interventions based, for example, on population size, in the interest of speed and workability. Utilizing economic–epidemiological modeling, we benchmark the performance of *ad hoc* allocation rules of scarce vaccines and drugs by comparing them to the rules for a vaccine and for a drug treatment that minimize the economic damages and expenditures over time, including a penalty cost representing the social costs of deviating from an *ad hoc* allocation. Under different levels of vaccine and drug scarcity, 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. Because drugs and vaccines attack different points in the disease pathology, the benefits from deviating from the *ad hoc* rule differ. For drug treatment, optimal policies often allocate all available treatments to one jurisdiction for a period of time, while *ad hoc* rules act to spread out treatments across jurisdictions. For vaccines, the benefits from deviating are especially high when immunity is permanent, when there is compliance to travel restrictions, and when the supply of vaccine is low. 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* rules, 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) and participants of the internal seminars at *Université du Québec à Montréal* (UQAM).

‡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, Knoxville, 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

While researchers race against time to develop vaccines and drug treatments against coronavirus disease 2019 (COVID-19), policymakers around the globe are focusing on how to allocate the limited supplies. Most of the scientific literature on allocation has focused on demographic considerations within one jurisdiction (Buckner et al., 2020; Emanuel et al., 2020a; Roope et al., 2020) or on a global scale (Emanuel et al., 2020a; Liu et al., 2020; Yamey et al., 2020). This prior work has made important contributions to the debate, but a missing piece in the allocation question is how to divide up limited quantities across jurisdictions that might have similar demographic but different epidemiological characteristics, because vaccine stockpiles may be managed at a federal level but distributed at more local scales (e.g. state). A recent report on the allocation of a future COVID-19 vaccine by the U.S. National Academies of Sciences, Engineering, and Medicine (2020) (NASEM) states that "[i]f the federal government were to provide states with an allotment of COVID-19 vaccine, in the interest of speed and workability, federal allocation to states could be conducted based on these jurisdictions' population size." Such a rule could also be deployed by states, provinces, or territories when deciding how to allocate within their boundaries.

In this paper, we explore the economic and epidemiological trade-offs associated with such fixed *ad hoc* allocation rules across jurisdictions by comparing them to the optimal rule for each jurisdiction conditional on the level of scarcity of the vaccine or drug. While most of the discussion revolves around 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, i.e. the virus that causes COVID-19, see Hu et al., 2020). 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. The baseline *ad hoc* allocation rules we consider are rules of thumb that favor "speed and workability." For vaccines, we follow the U.S. National Academies of Sciences, Engineering, and Medicine (2020) allocation based on the jurisdictions' population size. With respect to antiviral drugs, we base the *ad hoc* allocation on the jurisdictions' number of infected individuals.

In a world where two jurisdictions are identical in terms of population and infection levels, the *ad hoc* rules 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 vaccine or drug is developed. Mechanisms leading to heterogeneous disease burden include the timing of the outbreak, how preventive measures—such as mandatory face mask in public areas, shelter-in-place, travel restrictions, and social, or physical, distancing—vary from one area to another, and the degree to which local populations are compliant to preventive, or nonpharmaceutical measures (see Polyakova et al., 2020; Thomas et al., 2020, for more details on how SARS-CoV-2 prevalence varies across space).

Based on prior literature on spatial-dynamics of disease management, heterogeneity in disease burden is likely to lead to significant deviations in the optimal spatial allocation from the *ad hoc* rules (see Zaric and Brandeau, 2001, for example). Therefore, while compliance to preventive measures may seem independent from vaccine and drug allocation, they affect the pre-existing conditions (i.e. the conditions before the vaccine or drug is licensed) and the conditions under which the limited supplies will be allocated. For example, compliance to shelter-in-place and travel restrictions results in little to no leakage of the virus from one jurisdiction to another. When regions are non-interacting, Brandeau et al. (2003) 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 region and the initial level of infection. In the case of noncompliance to travel restrictions or populations are interacting, Rowthorn et al. (2009) show when there is no immunity that treatment should be preferentially directed towards the region that has the lower level of infection. In the case of COVID-19, whether noncompliance makes the *ad hoc* rule relatively more cost-effective is an open question.

The optimal allocations are conditional on the disease dynamics, compliance to travel restrictions, and the scarcity of the vaccines and drugs. In reality, the science is unresolved on whether immunity to SARS-CoV-2 is permanent or temporary, 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, 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 both drugs and vaccines over time to minimize the economic costs from damages, expenditures related to the pharmaceutical intervention, and a workability cost imposed on the planner for deviating from the *ad hoc* rules (i.e. the costs incurred because of the noncompliance to the *ad hoc* allocation). 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, Dangerfield et al., 2019). Second, we find that the vaccines and drugs should be optimally allocated over time depending on (i) if the jurisdiction has initially a lower or higher disease burden, (ii) if immunity is permanent (Zhou et al., 2014) or temporary, (iii) whether there is compliance to travel restrictions or not, and (iv) the amount of vaccine or drug available. We show that, compared to the *ad hoc* rule, the optimal allotment of drugs gives rise to extreme cases where all the drugs are given to one jurisdiction for a period of time. For the optimal allotment of vaccines, the benefits from deviating from the *ad hoc* rule are especially high when immunity is permanent, when there is compliance to travel restrictions, and when the supply is low. Third, we show that in general optimal rules are robust to uncertainty 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.

Our findings illustrate that allocating a future vaccine or drug based on an *ad hoc* allocation rule leads to an overutilization in jurisdictions where disease prevalence is higher, an underutilization in jurisdictions where disease prevalence is lower, and overall a higher number of cumulative cases. Whether these inefficiencies outweigh the "speed and workability" 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.

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 federal government) must decide when and how much of the scarce vaccines or drugs to allocate to two jurisdictions where disease burden is heterogeneous at the moment the vaccine or drug is licensed. 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.

2.1 Modelling Different Types of Intervention

There are three different types of interventions we consider: travel restrictions, drugs, 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 (see, e.g., Acemoglu et al., 2020; Alvarez et al., 2020, for examples of optimal lockdown policies). 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 effectively results in a discrete shift of the basic reproduction ratio, R_0 (see Appendix A for derivation).

We assume that at some given date either an antiviral drug or a vaccine is developed and licensed. We consider different scenarios of capacity constraints to investigate how different levels of limited supply may affect their optimal allocation. Antiviral drugs and vaccines affect different points in the disease pathology. The former shortens the infectious period while the latter reduces the pool of susceptible individuals by providing them with immunity from the disease. For simplicity, the amount of available vaccine or drug 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.

2.2 Model of Disease Transmission

We use a frequency-dependent (Begon et al., 2002) 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/administrative regions); each jurisdiction contains a population of N_i individuals (see Figure 1). We also consider scenarios where immunity is temporary (i.e. lasts 6 months, for more details see Edridge et al., 2020), thus also using an SEIR–Susceptible (SEIRS) model (for COVID-19 applications see, e.g., Bertozzi et al., 2020; Bjørnstad et al., 2020a,b; Hou et al., 2020; Pandey et al., 2020; Peng et al., 2020; Prem et al., 2020; Radulescu and Cavanagh, 2020; Roda et al., 2020; Stutt et al., 2020; Yang et al., 2020). In such scenarios, the R_i recovered individuals are immune for a mean period of $\frac{1}{\omega}$ months; in the case of permanent immunity ω is not defined.

In each jurisdiction i, the S_i susceptible individuals are in contact with the I_i infected individuals of their own jurisdiction at 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 abstract from other potential sources of heterogeneity across jurisdictions, we assume that the contact rate is identical across jurisdictions, meaning that $\beta_{11} = \beta_{22} = \beta_{ii}$ and $\beta_{12} = \beta_{21} = \beta_{ij}$. We are assuming there is no permanent migration of individuals from one jurisdiction to another (see for instance Burton et al., 2012, and see Chen et al., 2020 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 drugs or vaccines as non-proportional controls, i.e. available in a constant amount each month (Barrett and Hoel, 2007; Buckner et al., 2020; Goldman and Lightwood, 2002; Rowthorn et al., 2009). The change in susceptible individuals is

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

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. 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 revaccinated 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

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

Exposed individuals eventually become infectious for a mean period of time of $\frac{1}{\gamma+\varphi}$ and in turn can infect susceptible individuals. Infected individuals either recover naturally from the disease at a rate of γ, die from complications related to infection at a disease induced mortality rate of φ, or recover from the disease by being treated with antiviral drugs. The growth of the infected individuals is

$$\dot{I}_i = \sigma E_i - \gamma I_i - \varphi I_i - q_D u_{D_i} \tag{3}$$

where u_{D_i} represents the number of individuals being treated via antiviral drugs in a given time period (i.e. a month) in Jurisdiction i, and q_D represents the effectiveness of the drug.

The recovered population R_i includes individuals that recover naturally from the disease at a rate of γ, the individuals that are successfully vaccinated every month ($q_V u_{V_i}$), and the individuals that are successfully treated via drug every month ($q_D u_{D_i}$); if immunity is temporary ($\omega > 0$), a fraction of the recovered will leave this compartment. The number of recovered individuals in Jurisdiction i thus changes according to

$$\dot{R}_i = \gamma I_i + q_V u_{V_i} + q_D u_{D_i} - \omega R_i. \tag{4}$$

Note that 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

$$\dot{N}_i = -\varphi I_i. \tag{5}$$

We have omitted natural births and deaths due to the short time frame of our model. We also assume the population is closed, or that there is no exogenous importation of infected individuals (see for example Sivaraman et al., 2020, for an example applied to COVID-19) meant to mimic the fact that international travel has reduced dramatically

since the start of the pandemic (World Tourism Organization, 2020). See Appendix A for more details about the parameterization of the epidemiological model.

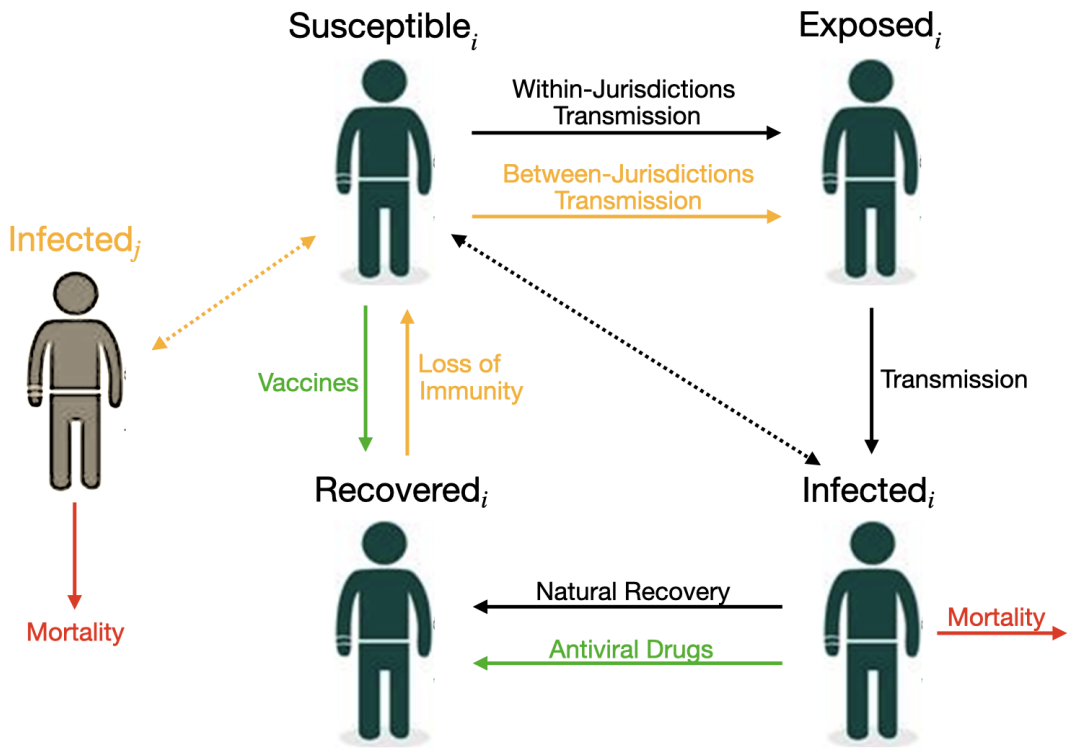

Figure 1: This schematic shows the model interventions and disease transmission pathways for our model of COVID-19. The full lines represent the transition between, or out of, compartments while the dotted lines represent contact between susceptible and infected individuals. Black lines represent situations 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 lines represent the pharmaceutical interventions and the red line represents mortality.

2.3 Modelling *Ad Hoc* Allocations

We model *ad hoc* allocation rules that favor "speed and workability." For vaccines, we follow the NASEM approach and impose that the allocation is based on relative population sizes. Specifically, the rule for Jurisdiction i is that

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

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

For drugs, the *ad hoc* rule is where the quantity of treatment allocated is proportional to the number of infected individuals. Specifically, the rule for Jurisdiction i is that

$$u_{D_i} \leq \left(\frac{I_i}{I_1 + I_2}\right)\bar{u}_D \tag{7}$$

where $\bar{u}_D$ is the limited amount of drug available for both jurisdictions. In general, as the level of infected individuals in the populations approaches zero, this formulation could lead to numerical instability. Given that this outcome is unlikely for COVID-19 and that we are considering very short time horizons in this analysis (months), we utilize this formulation.

In the *ad hoc* scenarios, we model both rules as an inequality because towards the end of the horizon after periods of vaccinations (or drug treatments), the level of susceptible (or infected) in the population may be such that the limited supply of vaccines (or drugs) is not an issue. Other *ad hoc* rules are possible, such as, allocate all to the largest or smallest population (Dangerfield et al., 2019), but we concentrate on the ones currently being advocated for by NASEM.

2.4 Model of Economic Costs

The model of economic costs include damages related to morbidity and deaths, costs spent on the drugs or 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, and the marginal value of damages is assumed to be constant over time and given by the value of a statistical life that the U.S. Environmental Protection Agency (EPA) 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 the World Health Organization (WHO) has not yet published disability values associated with COVID-19, following the literature (see for instance Nurchis et al., 2020), we use the disability value associated with lower respiratory tract infections. The damage function for Jurisdiction i is

$$c_i(I_i) = (w + \varphi)cI_i \tag{8}$$

where c is the damage parameter associated infectious individuals (i.e., the value of a statistical life; henceforth VSL).

We model a scenario where the central planner is focused on the allocation where the costs for the development of an effective drug or vaccine have already been incurred. This implies that these costs are sunk and do not affect the decision of the central planning agency. We model the cost of the pharmaceutical intervention as linear, where the cost parameter represents the cost of treating one individual. The cost function of the pharmaceutical intervention is denoted $c_{K_i}(\cdot)$, with $K = D, V$ and $i = 1, 2$, and where the arguments of the function depend on whether a vaccine or a drug is developed and becomes available. We assume that the cost of the pharmaceutical intervention is additively separable across jurisdictions such that we denote the cost of treating u_{K_i} individuals,

with $u_{K_i} = u_{V_i}, u_{D_i}$, as

$$c_{K_i}(u_{K_i}) = c_K u_{K_i} \text{ for } K = D, V \text{ and for } i = 1, 2, \tag{9}$$

where c_K represents the cost of treating one individual via drug ($K = D$) or vaccine ($K = V$). Calibration of the cost parameters are based on current expectations of the costs of drugs and vaccine prices (see Appendix A for more details about the parameterization of the economic model).

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 Ryan et al., 2017). For drugs, the workability cost function is:

$$c_A(u_{D_1}, u_{D_2}, I_1, I_2) = c_A\left(\left(\frac{I_2}{I_1 + I_2}\right)u_{D_1} - \left(\frac{I_1}{I_1 + I_2}\right)u_{D_2}\right)^2 \tag{10}$$

while for vaccines, the workability cost function is:

$$c_A(u_{V_1}, u_{V_2}, N_1, N_2) = c_A\left(\left(\frac{N_2}{N_1 + N_2}\right)u_{V_1} - \left(\frac{N_1}{N_1 + N_2}\right)u_{V_2}\right)^2 \tag{11}$$

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* rules *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 A13 for a sensitivity analysis of our results to the workability cost parameter when the intervention is an antiviral drug and see Figure A25 for a sensitivity analysis of our results to the workability cost parameter when the intervention is a vaccine). Therefore, when we do find deviations, we need to consider that these include this workability cost and if these costs do not exist, 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 and drugs) 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 Rowthorn et al., 2009). The objective function is the net present value of damages, expenditures related to the pharmaceutical intervention, and the

workability cost over an exogenously determined planning horizon (4 months). Specifically, the planner's objective is:

$$\min_{u_{K_1},u_{K_2}} \int_0^T e^{-rt}\Big\{c_1(I_1) + c_2(I_2) + c_{K_1}(u_{K_1}) + c_{K_2}(u_{K_2}) + c_A(\cdot)\Big\}\mathrm{d}t \tag{12}$$

where $K = D$ or $K = V$, and r is monthly discount rate. The planner solves Eq. (12) over a fixed time interval, T, subject to Eqs. (1), (2), (3), (4), (5), along with constraints on availability of the drug or vaccine ($u_{D_1} + u_{D_2} \leq \bar{u}_D$ or $u_{V_1} + u_{V_2} \leq \bar{u}_V$), non-negativity conditions, physical constraints on controls, initial disease burdens in each jurisdictions, and free endpoints. In the *ad hoc* scenarios, we also impose Eq. (6) for vaccines, and Eq. (7) for drugs. We consider the pharmaceutical interventions separately, that is, we consider the allocation of a vaccine where an antiviral drug does not exist and vice versa. In ongoing work, we consider the joint allocation question when both pharmaceutical interventions exist simultaneously.

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 vaccine or drug is licensed) is calibrated using the epidemiological model (Eqs.(1), (2), (3), (4), (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.

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 for each state variable that requires the product of the state variable 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 and antiviral drug over time, we numerically solve the optimal control problem across five different scenarios: no controls, optimal drug allocation, *ad hoc* drug allocation, optimal vaccine

allocation, and *ad hoc* vaccine allocation. We investigate how to allocate vaccines and drugs 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 or drug 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 Castonguay et al., 2020; Kling et al., 2016, for other applications, and see Appendix B for more details on this technique).

We present the results for our preferred specification of the parameters (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 pharmaceutical intervention is a drug or a vaccine, 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 or drug in a given time period (i.e. a month; $\bar{u}_D$ for drugs and $\bar{u}_V$ for vaccines) is based on a certain percentage (5%, 10%, or 15%) of the maximum level of infections in an uncontrolled outbreak (for drugs), and on a certain percentage (5%, 10%, or 15%) of the total population size (for vaccines). We focus our analysis on the period of time when the scarcity of the vaccine or drug constraint is binding, as once the constraint relaxes the allocation question becomes moot.

3.1 Allocation of Drugs

Regardless of whether the populations are compliant to travel restrictions, drugs should preferentially be directed towards the region that has the lower level of infection, when immunity is permanent (see Figure 2 Panel A for when there is compliance to travel restrictions). A similar counterintuitive result was found by Rowthorn et al. (2009). When immunity is temporary, however, the allocation of drugs initially favors the more infected state before switching back to the initially less infected state (see Figure 2 Panel B for when there is compliance to travel restrictions). This result concurs with the findings of Mbah and Gilligan (2011) which states that treatment should be allocated in a way that equalizes infection levels across jurisdictions as quickly as possible. When populations are non-compliant with travel restrictions, the mixing of the population leads to the "Rowthorn et al. (2009)" dominating the "Mbah and Gilligan (2011)" effect at higher levels of drug scarcity. Specifically, we find that the initially less infected state is favored by the optimal allocation (see Figure 3; the optimal allocation of drugs favors State 2 in Panel A, but it favors State 1 in Panel B). Noncompliance to travel restrictions does not affect qualitatively the allocation when there is permanent immunity (see Figure A1; Panel A is qualitatively similar to Panel B).

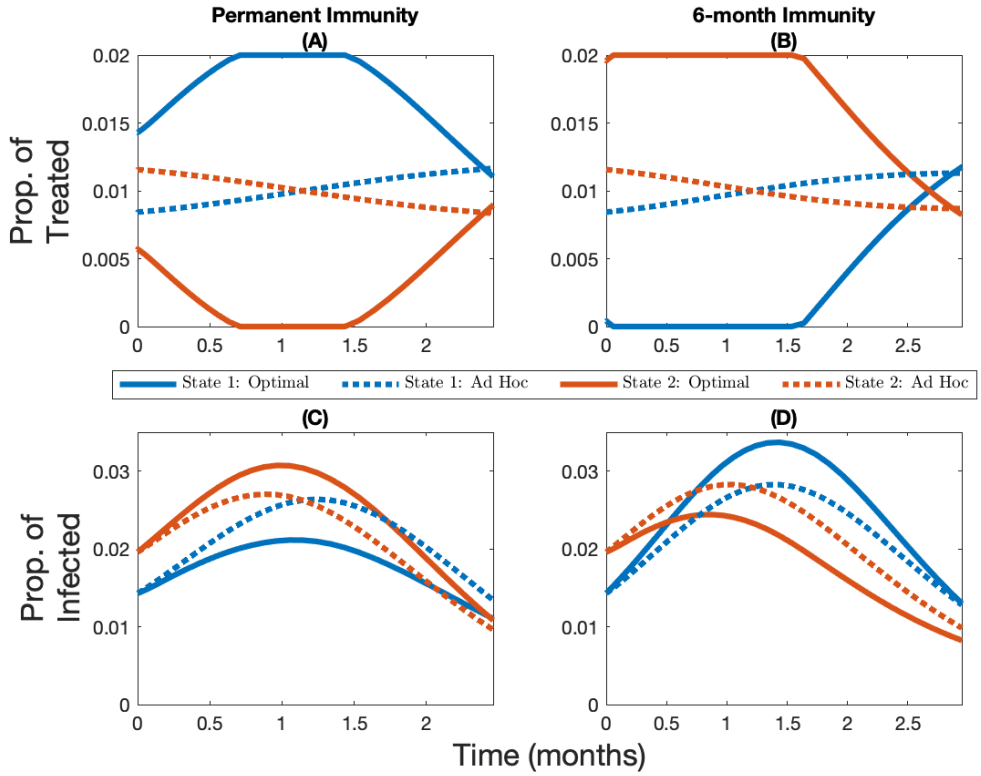

Figure 2: **Permanent vs temporary immunity with compliance to travel restrictions**. Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially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

Our findings imply that policy-makers governing in jurisdictions that comply with travel restrictions should change their allocation of drug treatment if immunity to COVID-19 turns out to be temporary (see Figure 2). When there is noncompliance to travel restrictions, the mixing of the populations mitigates the impact of the system's structural heterogeneity in terms of initial disease burden, and the effect of the temporary immunity has a small impact when drug capacity is low (see for example Figure A6 for when drug capacity is 5%). At higher drug supplies, temporary immunity does not qualitatively affect the result (e.g. 10%, see Figure A7).

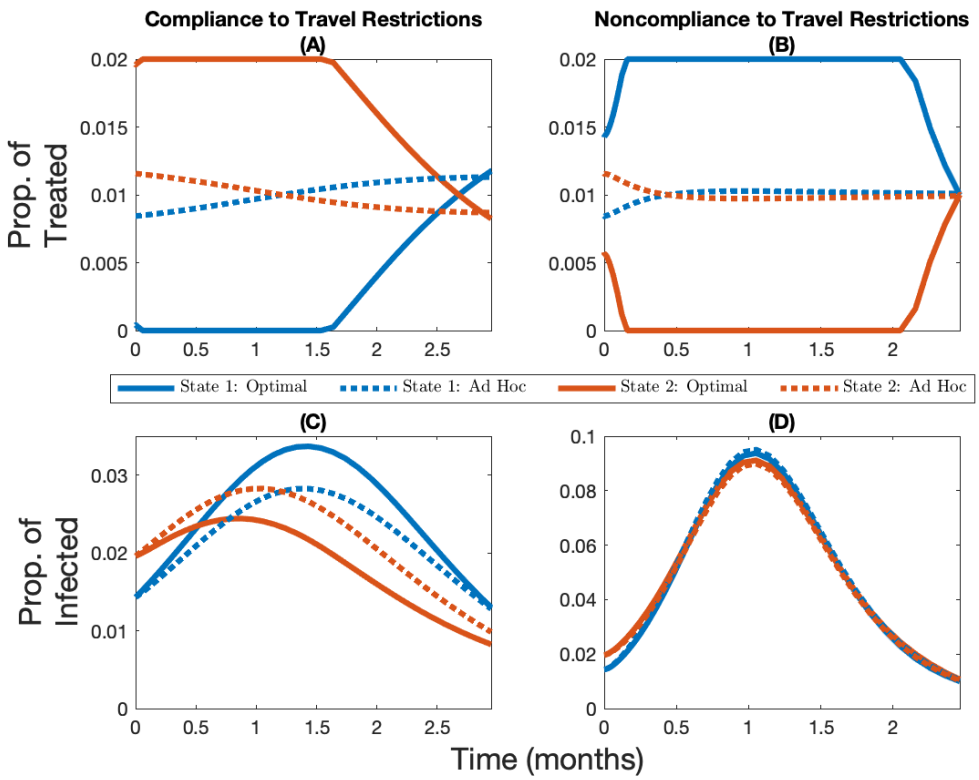

Figure 3: **Temporary immunity with and without compliance to travel restrictions**. Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially 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 pharmaceutical intervention is an antiviral drug, the drug capacity constraint is 10%, and immunity lasts six months.

Over the range we consider, an increase in the capacity constraint has a minimal impact on the drug allocation. Only when there is compliance to travel restrictions and a sufficiently high capacity does the optimal allocation move off of extreme cases where all the allotment is given to one state for a period of time (see Figure A2 and Figure A4, Panel C). In all of the other cases, there is a period of time where all the allotment of drugs goes optimally to one jurisdiction.

In all scenarios we consider, one state will have lower cumulative damages and one state will have higher cumulative damages as a result of the optimal allocation (see Figure 4 panels A, B, C, and D for the relative difference and panels E, F, G, and H for the absolute difference; one full line is below its dotted line while one full line is above its dotted line). Compared to the *ad hoc* allocation, the state with initially more infected individuals is worse off as a result of the optimal allocation in all cases except the one where immunity is temporary and populations comply with travel restrictions (this is due to the effect identified by Mbah and Gilligan, 2011; see Figure 4). Importantly, whether the decision-maker chooses to employ the *ad hoc* rule or the optimal rule, the effect of drugs on cumulative disease burden is considerably smaller when populations are noncompliant with the travel restrictions put in place;

this shows the importance of compliance to such nonpharmaceutical interventions (see Figure 4, and also figrues A8 and A9 in the appendix).

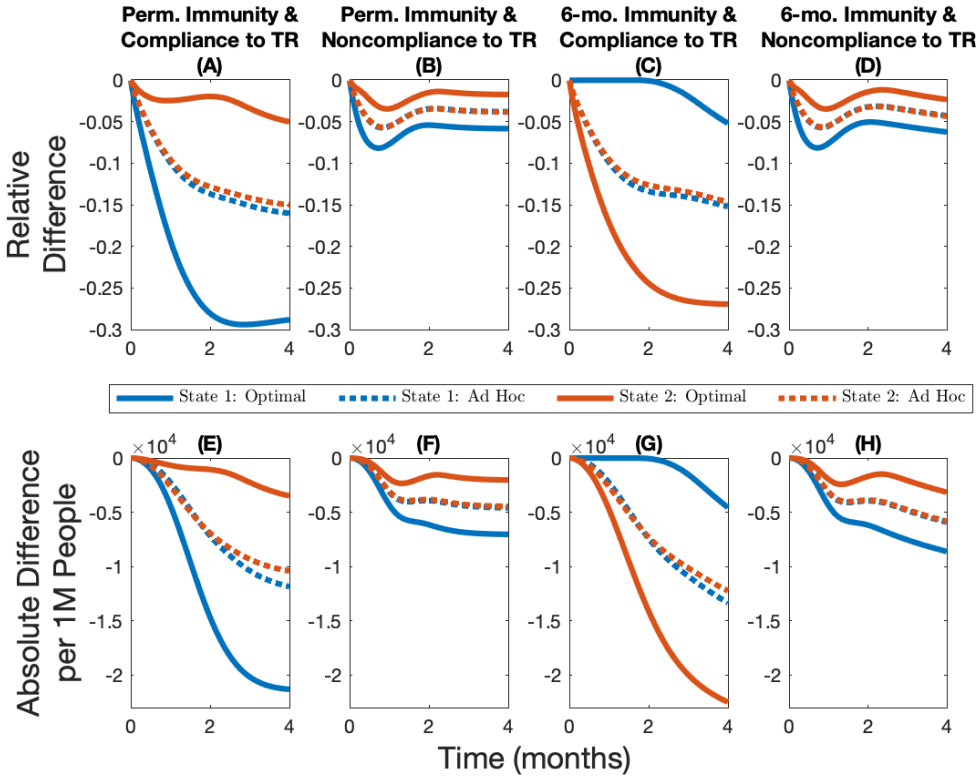

Figure 4: **Epidemiological outcomes under different scenarios with antiviral drug.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-drug case for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

Two key parameters in our analysis are the scale of the workability cost (c_A in Eq. 10) and level of drug effectiveness. While imposing the *ad hoc* rules *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 drug allocation over a range of values. We find that the deviation from the *ad hoc* allocation is stable for a wide range of the workablity cost parameter (Figure A13, panel A, B, C, and D), and that the optimal allocation will converge towards the *ad hoc* when the workability cost is in the neighborhood of the VSL (c in Eq. 8 and Figure A13 black line represents the VSL).

The base case effectiveness of the antiviral drugs is set based on a low "clinically meaningful" value (for more details see Appendix A). To better understand the impact of this key parameter on our results, we investigate a range of values. Overall, the deviation between the optimal and ad hoc is relatively stable, as evident by the scale of changes (in blue; Figure A14 Panels A-D). When there is non-compliance to travel restrictions, we find that higher levels of effectiveness increase the variance in the deviation between the optimal and *ad hoc* (in blue; Figure A14 panels

B and D). However, the relationship between optimal and *ad hoc* is more complicated when there is compliance, which recall corresponds to a lower R_o. Here we find a non-linear response (in blue; Figure A14 Panels A and C). In all cases, higher effectiveness means a larger difference in terms of cumulative cases between the optimal and *ad hoc* (in red Figure A14 Panels A-D).

3.2 Allocation of Vaccines

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 of the optimal allocation and the amplitude of the deviations from the *ad hoc* rule (see Figure 5 for when immunity is permanent, and see Figure A15 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 nonpharmaceutical interventions.

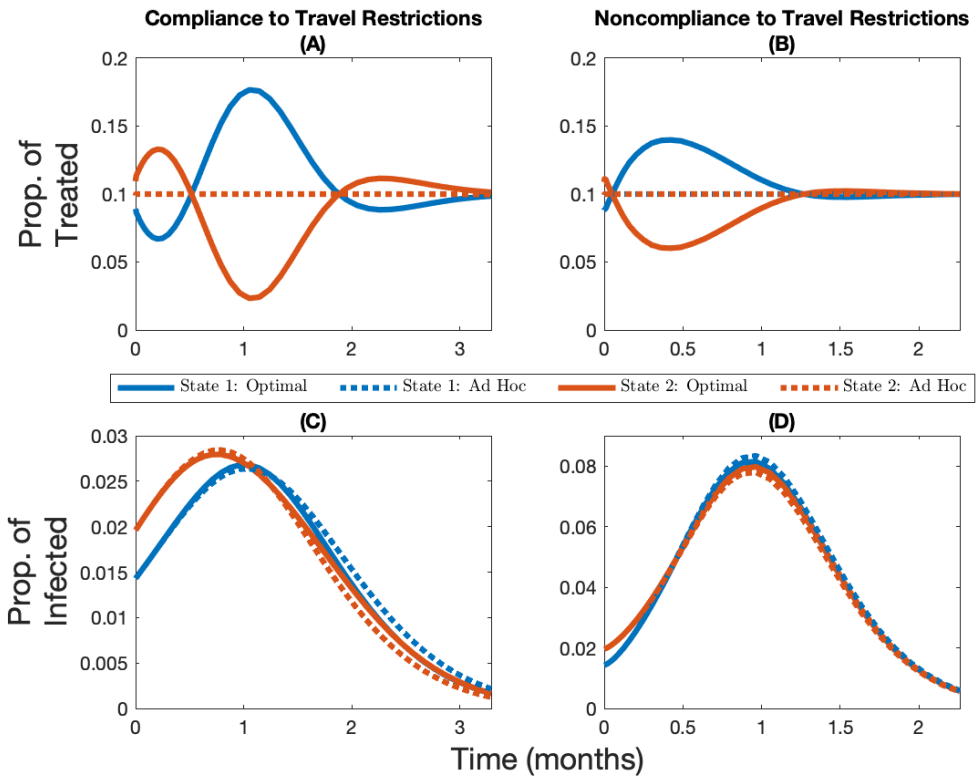

Figure 5: **Vaccine allocation with and without compliance to travel restrictions**. Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially 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 pharmaceutical intervention is a vaccine, the vaccine capacity constraint is 10%, and immunity is permanent.

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 6 Panel A for when immunity is permanent and Figure A18 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 6 panels B and C for when immunity is permanent and Figure A18 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 of permanent, a higher vaccine capacity implies that relatively more of the supply will be given to the more infected state at the beginning of the time horizon (see figures 6 and A16 for the case where immunity is permanent; see figures A17 and A18 for the case where immunity is temporary).

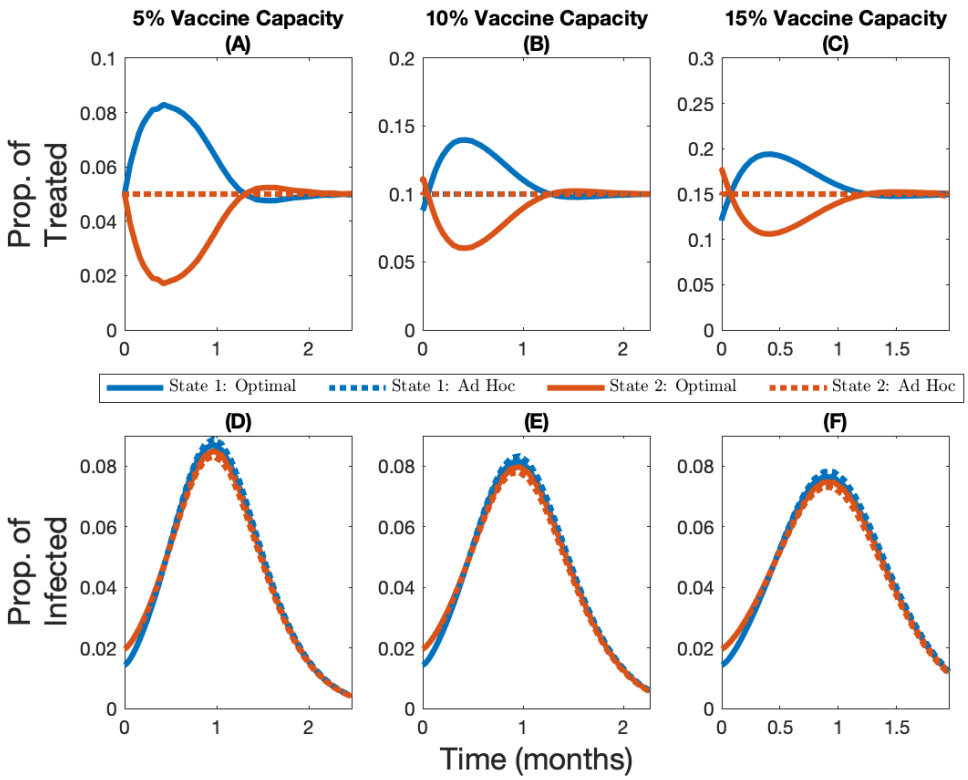

Figure 6: **Vaccine allocations under different levels of scarcity without compliance to travel restrictions**. Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened 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 the pharmaceutical intervention is a vaccine, 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 A19). 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, the allocation equalizes the current infection levels across jurisdictions (Figure 5 Panel A). As the vaccine capacity increases, however, the *ad hoc* allocation rule performs better that in turn decreases the amplitude of the optimal deviation (see Figure A16 panels A, B and C, or Figure A17 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 7). Hence, the optimal allocation makes current infection level more equal, while *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 7 with vaccine capacity is 10%, and see figures A21 and A22 with vaccine capacity is 5% and 15%, respectively).

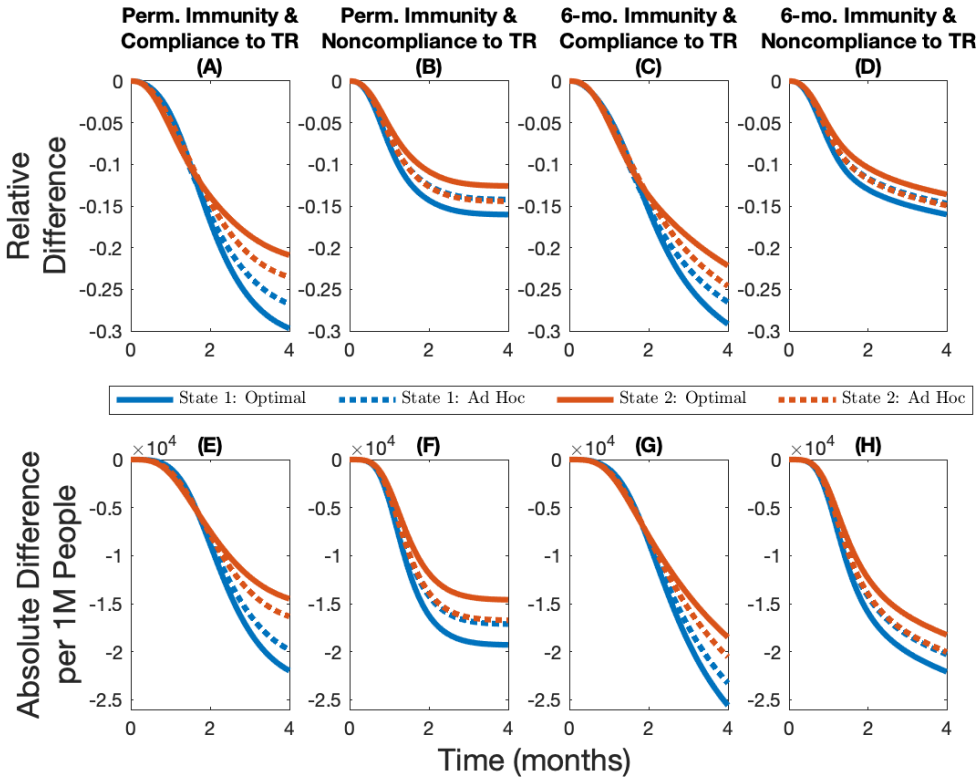

Figure 7: **Epidemiological outcomes under different scenarios with vaccines.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-vaccine case for different immunity–travel restrictions scenarios and for when vaccine capacity is 10%.

Similar to the drug case, we investigate the impact of the scale of the workability cost (c_A in Eq. 11) and level of vaccine effectiveness (see Appendix C for more details). 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 A25 panels A, B, C, and D). Specifically, we find that when the cost is in the neighborhood of the VSL (c in Eq. 8 and Figure A25 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 (Ohmit et al., 2014, see Appendix A for more details). Recent evidence from the COVID-19 vaccines suggest that effectiveness could be considerably higher. 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 A26 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 A26 panels A, B, C, and D).

3.3 Robustness of Spatial Allocations

Given significant uncertainty associated with the duration of immunity (i.e. if it is permanent or temporary), or to what extent populations comply with travel restrictions, we compare the robustness of the optimal spatial allocation to the *ad hoc*. By definition, the optimal allocation minimizes the net present value of the economic damages and expenditures related to both the intervention and deviations from the *ad hoc* allocation, and thus cannot do worse on these dimensions 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 8). 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 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 of vaccines and include the drug and other scarcity cases in Appendix C.

For vaccine allocations, 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 8). There are more nuanced trade-offs, however (e.g. compare position of the stars across the panels in Figure 8). 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 8 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 A23 and A24 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 A24).

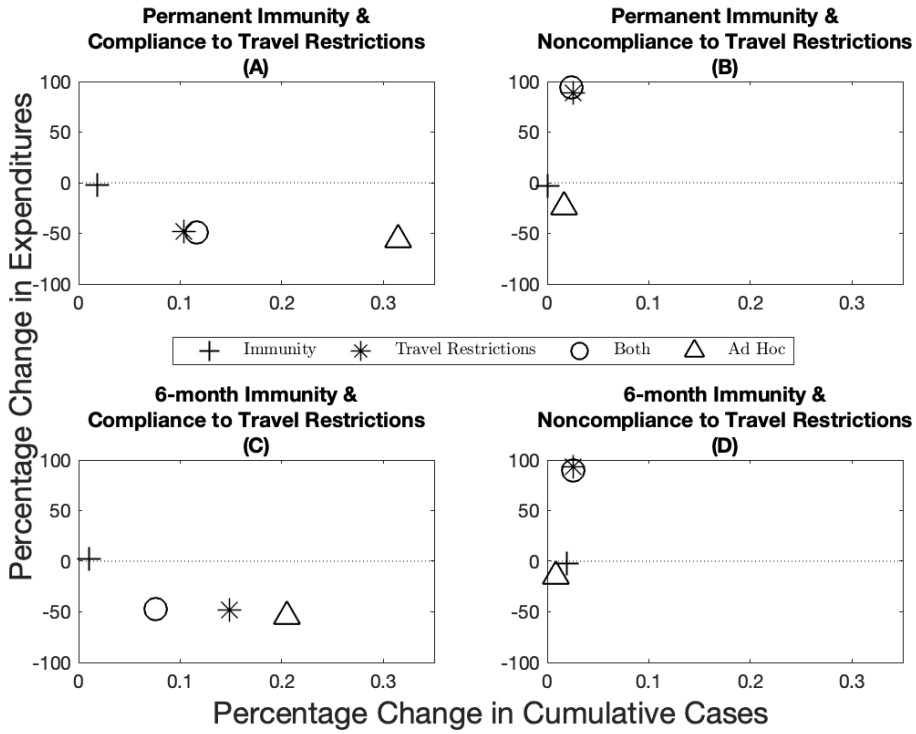

Figure 8: **Robustness of epidemiological and economic outcomes under different scenarios with vaccines.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine 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.

For drug allocations, the optimal allocations are generally not robust to incomplete information and perform worse than the *ad hoc* allocations (see figures A10, A11, and A12 for when drug capacity is 5%, 10%, and 15% respectively). Recall that the *ad hoc* in this case is based off of level of infected individuals in each jurisdiction, which changes over time, and captures the disease burden well. We find that when populations are compliant to travel restrictions

and drug capacity increases, the optimal allocation converges towards the *ad hoc* allocation (see Figures A2 and A4). This implies that the *ad hoc* performs better as drug capacity increases. But this result only holds when populations comply to travel restrictions.

When populations do not comply to travel restrictions yet drugs are allocated as if there is compliance, the optimal drug allocation gets less robust as supply increases (figures A10-A12 Panel C). When populations are noncompliant to travel restrictions, the transmission of the disease is considerably higher, which increases the basic reproduction ratio. In this case regardless of the allocation rule, the effect that the drug treatment has on the infection rate is relatively smaller (i.e., treating 1% of the population when 3% is infected does not have the same overall impact as treating 1% of the population when 10% is infected). Hence, the gains of switching from the *ad hoc* to the optimal allocation are relatively smaller when there is noncompliance to travel restrictions.

4 Conclusion

Recent studies have discussed how a future vaccine against the coronavirus disease (COVID-19) should be allocated on a global scale (see for instance Emanuel et al., 2020b; Liu et al., 2020; World Health Organization and others, 2020; Yamey et al., 2020) and within a geographical area (see for instance Buckner et al., 2020; Emanuel et al., 2020a; Roope et al., 2020). Building off the spatial-dynamic literature in epidemiology, we contribute to this body of work by addressing the question of distributing a scarce COVID-19 vaccine and antiviral drug across smaller geographic areas, such as counties, regions, or states. The U.S. National Academies of Sciences, Engineering, and Medicine (2020) recommends to allocate a future vaccine against COVID-19 based on the jurisdictions' population size. In this paper, we show the potential economic and public health benefits of deviating from an *ad hoc* allocation rule, which in turn provides policymakers with information on the trade-offs involved with different allocations. 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, and the capacity constraint 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. (2009) 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 (Duijzer et al., 2018). We find that temporary immunity can lead to the opposite allocation, which is consistent with the findings in Mbah and Gilligan (2011).

We also show the value of complying to a travel restriction, as compliance leads to lower cumulative damages across both jurisdictions, regardless of whether the pharmaceutical intervention is a vaccine or a drug, and 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 (Verelst et al., 2020) and a corresponding varying death rate due to scarce intensive care unit beds (Acemoglu et al., 2020), and other second-order problems such as consumption losses (Andersen et al., 2020; Baker et al., 2020b; Coibion et al., 2020), excess mortality (Vestergaard et al., 2020), and psychological distress (Pfefferbaum and North, 2020) 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 federal government) to deviate from this rule of thumb; this result holds in all scenarios we considered in our analysis. We considered *ad hoc* allocation rules that favor "speed and workability" (National Academies of Sciences, Engineering, and Medicine, 2020). Other allocation rules are possible. For instance, in our model 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 resources to their jurisdictions.

Extrapolating our results to the entire U.S. suggests that allocating a vaccine based on the *ad hoc* allocation rules in this paper can have serious public health consequences. How many additional cases depends on several factors including epidemiological (i.e. length of immunity), behavioral (i.e. compliance to travel restrictions), and logistical (i.e. vaccine capacity) factors. In the United States alone and with 10 percent capacity, the increase in the number of cases due to an allocation of a scarce COVID-19 vaccine based on the relative population size of the states could imply as little as 28,000 additional cases, but according to our model this number could be as high as 1.03 million additional cases. Fortunately, additional vaccine capacity in the range considered in the paper improves the relative performance of the *ad hoc* allocation when there is compliance to travel restrictions, but at the same time, the performance of the *ad hoc* allocation when there is noncompliance to travel restrictions is worsened. For instance, when vaccine capacity is 5%, the range goes from 21,000 to 1.05 million additional cases, and when vaccine capacity is 15%, the range goes from 34,000 to 950,000 additional cases. For antiviral drugs, a similar troubling result holds. Allocating a low supply of a future antiviral drug based the relative number of infected individuals could imply as few as $10,000$ additional cases but this number could increase to 2.4 million. Unfortunately, this problem does not go away as drug supply increases, as a higher supply of antiviral drug would increase this range to between 37,000 and 2.9 million additional cases.

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 our population is assumed to be homogeneous and the populations across jurisdictions are assumed to have identical characteristics. In practice, however, the virus disproportionately affects elderly people (Verity et al., 2020) and people with pre-existing conditions (Ssentongo et al., 2020). Also, the risk of infection is highly occupational dependent (Baker et al., 2020a). 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.

References

Acemoglu, D., Chernozhukov, V., Werning, I., and Whinston, M. D. (2020). Optimal targeted lockdowns in a multi-group SIR model. *NBER Working Paper*, 27102.

Alvarez, F. E., Argente, D., and Lippi, F. (2020). A simple planning problem for COVID-19 lockdown. *NBER Working Paper*, 26981.

Andersen, A. L., Hansen, E. T., Johannesen, N., and Sheridan, A. (2020). Consumer responses to the COVID-19 crisis: Evidence from bank account transaction data. *Available at SSRN 3609814*.

Baker, M. G., Peckham, T. K., and Seixas, N. S. (2020a). Estimating the burden of united states workers exposed to infection or disease: a key factor in containing risk of COVID-19 infection. *PLoS One*, 15(4):e0232452.

Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., and Yannelis, C. (2020b). How does household spending respond to an epidemic? consumption during the 2020 COVID-19 pandemic. Technical report, National Bureau of Economic Research.

Barrett, S. and Hoel, M. (2007). Optimal disease eradication. *Environment and Development Economics*, pages 627–652.

Begon, M., Bennett, M., Bowers, R. G., French, N. P., Hazel, S., and Turner, J. (2002). A clarification of transmission terms in host-microparasite models: numbers, densities and areas. *Epidemiology & Infection*, 129(1):147–153.

Bertozzi, A. L., Franco, E., Mohler, G., Short, M. B., and Sledge, D. (2020). The challenges of modeling and forecasting the spread of COVID-19. *arXiv preprint arXiv:2004.04741*.

Bjørnstad, O., Shea, K., Krzywinski, M., and Altman, N. (2020a). The SEIRS model for infectious disease dynamics. *Nature Methods*, 17(6):557–558.

Bjørnstad, O. N., Shea, K., Krzywinski, M., and Altman, N. (2020b). Modeling infectious epidemics. *Nature methods*, 17:455–456.

Brandeau, M. L., Zaric, G. S., and Richter, A. (2003). Resource allocation for control of infectious diseases in multiple independent populations: beyond cost-effectiveness analysis. *Journal of health economics*, 22(4):575–598.

Buckner, J. H., Chowell, G., and Springborn, M. R. (2020). Optimal dynamic prioritization of scarce COVID-19 vaccines. *medRxiv*.

Burton, J., Billings, L., Cummings, D. A., and Schwartz, I. B. (2012). Disease persistence in epidemiological models: the interplay between vaccination and migration. *Mathematical biosciences*, 239(1):91–96.

Castonguay, F. M., Sokolow, S. H., De Leo, G. A., and Sanchirico, J. N. (2020). Cost-effectiveness of combining drug and environmental treatments for environmentally transmitted diseases. *Proceedings of the Royal Society B*, 287(1933):20200966.

Chen, M., Li, M., Hao, Y., Liu, Z., Hu, L., and Wang, L. (2020). The introduction of population migration to seiar for COVID-19 epidemic modeling with an efficient intervention strategy. *Information Fusion*, 64:252–258.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). The cost of the COVID-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. Technical report, National Bureau of Economic Research.

Dangerfield, C. E., Vyska, M., and Gilligan, C. A. (2019). Resource allocation for epidemic control across multiple sub-populations. *Bulletin of mathematical biology*, 81(6):1731–1759.

Duijzer, L. E., van Jaarsveld, W. L., Wallinga, J., and Dekker, R. (2018). Dose-optimal vaccine allocation over multiple populations. *Production and Operations Management*, 27(1):143–159.

Edridge, A. W., Kaczorowska, J. M., Hoste, A. C., Bakker, M., Klein, M., Jebbink, M. F., Matser, A., Kinsella, C., Rueda, P., Prins, M., et al. (2020). Coronavirus protective immunity is short-lasting. *MedRxiv*.

Emanuel, E., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., and Phillips, J. (2020a). Fair allocation of scarce medical resources in the time of COVID-19. *The New England Journal of Medicine*.

Emanuel, E. J., Persad, G., Kern, A., Buchanan, A., Fabre, C., Halliday, D., Heath, J., Herzog, L., Leland, R., Lemango, E. T., et al. (2020b). An ethical framework for global vaccine allocation. *Science (New York, NY)*, page eabe2803.

Goldman, S. M. and Lightwood, J. (2002). Cost optimization in the sis model of infectious disease with treatment. *The BE Journal of Economic Analysis & Policy*, 2(1).

Hou, C., Chen, J., Zhou, Y., Hua, L., Yuan, J., He, S., Guo, Y., Zhang, S., Jia, Q., Zhao, C., et al. (2020). 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*.

Hu, B., Guo, H., Zhou, P., and Shi, Z.-L. (2020). Characteristics of sars-cov-2 and covid-19. *Nature Reviews Microbiology*, pages 1–14.

Kling, D. M., Sanchirico, J. N., and Wilen, J. E. (2016). Bioeconomics of managed relocation. *Journal of the Association of Environmental and Resource Economists*, 3(4):1023–1059.

Liu, Y., Salwi, S., and Drolet, B. (2020). Multivalue ethical framework for fair global allocation of a covid-19 vaccine. *Journal of Medical Ethics*.

Mbah, M. L. N. and Gilligan, C. A. (2011). Resource allocation for epidemic control in metapopulations. *PLoS one*, 6(9):e24577.

National Academies of Sciences, Engineering, and Medicine (2020). Discussion draft of the preliminary framework for equitable allocation of COVID-19 vaccine. `https://doi.org/10.17226/25914`.

Nurchis, M. C., Pascucci, D., Sapienza, M., Villani, L., D'Ambrosio, F., Castrini, F., Specchia, M. L., Laurenti, P., and Damiani, G. (2020). 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*, 17(12):4233.

Ohmit, S. E., Thompson, M. G., Petrie, J. G., Thaker, S. N., Jackson, M. L., Belongia, E. A., Zimmerman, R. K., Gaglani, M., Lamerato, L., Spencer, S. M., et al. (2014). Influenza vaccine effectiveness in the 2011–2012 season: protection against each circulating virus and the effect of prior vaccination on estimates. *Clinical infectious diseases*, 58(3):319–327.

Pandey, G., Chaudhary, P., Gupta, R., and Pal, S. (2020). SEIR and regression model based COVID-19 outbreak predictions in India. *arXiv preprint arXiv:2004.00958*.

Peng, L., Yang, W., Zhang, D., Zhuge, C., and Hong, L. (2020). Epidemic analysis of COVID-19 in China by dynamical modeling. *arXiv preprint arXiv:2002.06563*.

Pfefferbaum, B. and North, C. S. (2020). Mental health and the covid-19 pandemic. *New England Journal of Medicine*.

Polyakova, M., Kocks, G., Udalova, V., and Finkelstein, A. (2020). 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*.

Prem, K., Liu, Y., Russell, T. W., Kucharski, A. J., Eggo, R. M., Davies, N., Flasche, S., Clifford, S., Pearson, C. A., Munday, J. D., et al. (2020). 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*.

Radulescu, A. and Cavanagh, K. (2020). Management strategies in a SEIR model of COVID 19 community spread. *arXiv preprint arXiv:2003.11150*.

Roda, W. C., Varughese, M. B., Han, D., and Li, M. Y. (2020). Why is it difficult to accurately predict the COVID-19 epidemic? *Infectious Disease Modelling*.

Roope, L. S., Buckell, J., Becker, F., Candio, P., Violato, M., Sindelar, J. L., Barnett, A., Duch, R., and Clarke, P. M. (2020). How should a safe and effective covid-19 vaccine be allocated? health economists need to be ready to take the baton. *PharmacoEconomics-Open*, pages 1–5.

Rowthorn, R. E., Laxminarayan, R., and Gilligan, C. A. (2009). Optimal control of epidemics in metapopulations. *Journal of the Royal Society Interface*, 6(41):1135–1144.

Ryan, D., Toews, C., Sanchirico, J. N., and Armsworth, P. R. (2017). Implications of policy adjustment costs for fisheries management. *Natural Resource Modeling*, 30(1):74–90.

Sivaraman, N. K., Gaur, M., Baijal, S., Rupesh, C. V., Muthiah, S. B., and Sheth, A. (2020). Exo-sir: An epidemiological model to analyze the impact of exogenous infection of COVID-19 in india. *arXiv preprint arXiv:2008.06335*.

Ssentongo, P., Ssentongo, A. E., Heilbrunn, E. S., Ba, D. M., and Chinchilli, V. M. (2020). Association of cardiovascular disease and 10 other pre-existing comorbidities with COVID-19 mortality: A systematic review and meta-analysis. *PloS one*, 15(8):e0238215.

Stutt, R. O., Retkute, R., Bradley, M., Gilligan, C. A., and Colvin, J. (2020). 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*, 476(2238):20200376.

Thomas, L. J., Huang, P., Yin, F., Luo, X. I., Almquist, Z. W., Hipp, J. R., and Butts, C. T. (2020). Spatial heterogeneity can lead to substantial local variations in COVID-19 timing and severity. *arXiv preprint arXiv:2005.09850*.

Verelst, F., Kuylen, E., and Beutels, P. (2020). Indications for healthcare surge capacity in european countries facing an exponential increase in coronavirus disease (COVID-19) cases, march 2020. *Eurosurveillance*, 25(13):2000323.

Verity, R., Okell, L. C., Dorigatti, I., Winskill, P., Whittaker, C., Imai, N., Cuomo-Dannenburg, G., Thompson, H., Walker, P. G., Fu, H., et al. (2020). Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet infectious diseases*.

Vestergaard, L. S., Nielsen, J., Richter, L., Schmid, D., Bustos, N., Braeye, T., Denissov, G., Veideman, T., Luomala, O., Möttönen, T., et al. (2020). Excess all-cause mortality during the COVID-19 pandemic in europe–preliminary pooled estimates from the euromomo network, march to april 2020. *Eurosurveillance*, 25(26):2001214.

World Health Organization and others (2020). Who sage values framework for the allocation and prioritization of COVID-19 vaccination, 14 september 2020. Technical report, World Health Organization.

World Tourism Organization (2020). *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).

Yamey, G., Schäferhoff, M., Hatchett, R., Pate, M., Zhao, F., and McDade, K. K. (2020). Ensuring global access to COVID-19 vaccines. *The Lancet*, 395(10234):1405–1406.

Yang, Z., Zeng, Z., Wang, K., Wong, S.-S., Liang, W., Zanin, M., Liu, P., Cao, X., Gao, Z., Mai, Z., et al. (2020). Modified SEIR and AI prediction of the epidemics trend of COVID-19 in China under public health interventions. *Journal of Thoracic Disease*, 12(3):165.

Zaric, G. S. and Brandeau, M. L. (2001). Optimal investment in a portfolio of HIV prevention programs. *Medical Decision Making*, 21(5):391–408.

Zhou, Y., Yang, K., Zhou, K., and Liang, Y. (2014). Optimal vaccination policies for an SIR model with limited resources. *Acta biotheoretica*, 62(2):171–181.

Appendices

Contents

A Parameterization **29**
A.1 Epidemiological Model . . . 29
A.2 Economic Model . . . 29
A.3 Parameter Levels . . . 30

B Optimization **31**
B.1 Boundary Conditions . . . 31
B.2 Nonnegativity and Upper-Bound Constraints . . . 32
B.3 Capacity Constraints of the Pharmaceutical Interventions . . . 32
B.4 Numerical Methods . . . 32

C Figures **33**
C.1 Drugs . . . 33
C.1.1 Compliance and Noncompliance to Travel Restrictions . . . 33
C.1.2 Drug Capacity Constraints when Immunity is Permanent . . . 34
C.1.3 Drug Capacity Constraints when Immunity is Temporary . . . 35
C.1.4 Permanent vs Temporary Immunity . . . 36
C.1.5 Cumulative Infection Levels . . . 37
C.1.6 Robustness of Optimal Allocations . . . 39
C.1.7 Sensitivity Analysis of Workability Cost . . . 41
C.1.8 Sensitivity Analysis of Drug Effectiveness . . . 43
C.2 Vaccines . . . 44
C.2.1 Compliance and Noncompliance to the Travel Restrictions . . . 44
C.2.2 Vaccine Capacity Constraints when Immunity is Permanent . . . 45
C.2.3 Vaccine Capacity Constraints when Immunity is Temporary . . . 46
C.2.4 Permanent vs Temporary Immunity . . . 47
C.2.5 Cumulative Infection Levels . . . 48
C.2.6 Robustness of Optimal Allocations . . . 50
C.2.7 Sensitivity Analysis of Workability Cost . . . 51
C.2.8 Sensitivity Analysis of Vaccine Effectiveness . . . 53

A Parameterization

A.1 Epidemiological Model

According to Diekmann et al. (1990), the basic reproduction ratio R_0 of any disease is given by the expected number of secondary infection caused a 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-state SEIR model in this paper—one needs to use the next-generation matrix and find its dominant eigenvalue to find the R_0 (Diekmann et al., 1990). 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 (Garchitorena et al., 2017); the next-generation matrix is equal to $-FV^{-1}$. In the model presented in this paper,

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

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 Diekmann et al., 2010). Given we assume that $\beta_{11} = \beta_{22} = \beta_{ii}$ and $\beta_{12} = \beta_{21} = \beta_{ij}$, 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. (2020) and using estimates of the effect of nonpharmaceutical interventions on the R_0 from Tian et al. (2020). We assume a mean recovery period $(\frac{1}{\gamma})$ of 5 days (Davies et al., 2020), and a case-fatality rate of 1.78% (adjusted for misreporting, Abdollahi et al., 2020) to calibrate the rate of disease induced mortality, φ. Parameters β_{ii} and β_{ij} are then calibrated assuming what Tian et al. (2020) 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 Schlosser et al., 2020); 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 (Davies et al., 2020).

A.2 Economic Model

To quantify damages, we use the value of statistical life recommended by the Environmental Protection Agency.[1] The disability weight[2] 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 what a future drug or vaccine could cost. For drugs our parameter value is based on the drug REGN-COV2—produced by the biotech company Regeneron—that President Trump took after his positive diagnosis to COVID-19.[4] For vaccines, numerous governments around the world, including the U.S. federal government, have contracted biotech companies producing

[1] See "What value of statistical life does EPA use?" from the U.S. Environmental Protection Agency (2020).

[2] According 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/`.

[3] For more details on how COVID-19's disability resembles lower respiratory tract infections, see Nurchis et al. (2020).

[4] The drug is expected to cost between \$1,500 and \$6,500 per patient; see: `https://www.cbsnews.com/news/what-is-regeneron-covid-antibody-cocktail-trump-covid-19/`

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 turns out to be approximately 20 U.S. dollars per dose, with two doses per individual; this is the value we chose in our analysis.[5]

The value of the workability cost[6] 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 John et al., 2019, 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}).[7]
β_{ij}	4.36	Transmission rate across states (month^{-1}).[7]
σ	10.14	Rate at which infected individuals become infectious (month^{-1}).[8]
γ	6.08	Rate of recovery (month^{-1}).[8]
ω	0.17	Rate at which immunity is lost (month^{-1}).[9]
φ	0.11	Rate of disease induced mortality (month^{-1}).[10]
w	0.13	Disability weight associated with the disease (unitless).[11]
q_D	0.65	Efficiency of drugs (proportion).[12]
q_V	0.65	Efficiency of vaccines (proportion).[13]
r	0.0013	Discount rate (month^{-1}).[14]
c_D	1,500	Cost of treating one individual via drugs (US Dollars).[15]
c_V	40	Cost of treating one individual via vaccine (US Dollars).[16]
c_A	10×10^3	Workability cost (US Dollars).[17]
c	10×10^6	Value of statistical life (US Dollars).[18]

Table A1: Parameter levels used in the numerical simulation.

[5]For 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`.

[6]Inspired by the paper of Ryan et al. (2017) where the authors show the implications of policy adjustment costs for fisheries management

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 (Tian et al., 2020) (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	0.9572	0.0102	0.0112	0.0213
Jurisdiction 2	1	0.9134	0.0199	0.0224	0.0441

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:

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

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

[7] Calibrated using a R_0 estimate from Li et al. (2020) and estimates of effects of nonpharmaceutical interventions from Tian et al. (2020); 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.

[8] Using estimates from Davies et al. (2020); this represents a 3-day latency period and a 5-day recovery period.

[9] Representing a 6-month immunity period in the scenarios where we assume immunity is not permanent; based on Edridge et al. (2020).

[10] Calibrated by using a case-fatility rate of 1.78% (adjusted for mis- and under-reporting; see Abdollahi et al., 2020).

[11] Representing the disability associated with severe lower respiratory tract infections because, to our knowledge, there are no official disability estimates associated with COVID-19; see Nurchis et al. (2020).

[12] According to the U.S. Department of Health and Human Services, Food and Drug Administration (FDA), Center for Biologics Evaluation and Research (CBER), and Center for Drug Evaluation and Research (CDER) (2019), there is no required effectiveness for a newly developed drug, however, what is necessary is that: "it is well established that the effect shown in the adequate and well-controlled clinical investigations, must be, in FDA's judgment, clinically meaningful." As a result, and to be consistent across both types of pharmaceutical interventions, we assume in the base case the effectiveness of the drug is 0.65.

[13] Following Buckner et al. (2020), we base this parameter value on the efficiency of the influenza vaccine (see Ohmit et al., 2014). Note that the U.S. Department of Health and Human Services, Food and Drug Administration (FDA), and Center for Biologics Evaluation and Research (CBER) (2020) requires that a future COVID-19 vaccine must have an effectiveness of at least 50%.

[14] Based on results from John et al. (2019) 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 $r = 0.0013$.

[15] Based on current expectations of the costs of REGN-COV2, the medicine produced by Regeneron that President Trump took following his COVID-19 diagnosis. According to information from U.S. Food and Drug Administration (FDA), CBS News, and Regeneron, the drug is expected to cost between $1,500 and $6,500 depending on the exact number of doses required by a patient; see: `https://www.cbsnews.com/news/what-is-regeneron-covid-antibody-cocktail-trump-covid-19/`.

[16] Assuming 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`.

[17] Value 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.

[18] Represents 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 (2020) uses: approximately $7.4 million ($2006) which is equivalent to approximately $9.54 million ($2020).

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:

$$0 \leq S_i \leq N_i \leq 1 \text{ for } i = 1, 2; \tag{A2a}$$
$$0 \leq E_i \leq N_i \leq 1 \text{ for } i = 1, 2; \tag{A2b}$$
$$0 \leq I_i \leq N_i \leq 1 \text{ for } i = 1, 2; \tag{A2c}$$
$$0 \leq R_i \leq N_i \leq 1 \text{ for } i = 1, 2; \tag{A2d}$$
$$S_i + E_i + I_i + R_i = N_i \leq 1 \text{ for } i = 1, 2. \tag{A2e}$$

Control variables are modelled as direct controls (see examples in Barrett and Hoel, 2007; Goldman and Lightwood, 2002; Rowthorn et al., 2009) and can be interpreted as a reduction in the number of infected individuals (for drugs) or susceptible individuals (for vaccines) in a given time period (i.e. a month). Formally, the constraints on the control variables are given by:

$$0 \leq u_{D_i} \leq I_i \text{ for } i = 1, 2; \tag{A3a}$$
$$0 \leq u_{V_i} \leq S_i \text{ for } i = 1, 2. \tag{A3b}$$

Because of a limited supply of drugs and 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 infected individuals than there are available drugs, or 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 limited supplies of a newly licensed vaccine or drug, before the supply has had a chance to ramp up. Hence, the control variables are also subject to

$$u_{D_1} + u_{D_2} \leq \bar{u}_D; \tag{A4a}$$
$$u_{V_1} + u_{V_2} \leq \bar{u}_V; \tag{A4b}$$

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

$$u_{D_i} \leq \left(\frac{I_i}{I_1 + I_2}\right) \bar{u}_D \qquad \text{for } i = 1, 2; \tag{A5a}$$

$$u_{V_i} \leq \left(\frac{N_i}{N_1 + N_2}\right) \bar{u}_V \quad \text{for } i = 1, 2. \tag{A5b}$$

As mentioned in the main paper, the total available quantity of vaccine or drug ($\bar{u}_D$ for drugs and $\bar{u}_V$ for vaccines) represents a certain percentage (5%, 10%, or 15%) of the maximum level of infections in an uncontrolled outbreak (for drugs), and on a certain percentage (5%, 10%, or 15%) of the total population size (for vaccines).

B.4 Numerical Methods

Pseudospectral collocation approximates the continuous time optimal control model with a constrained nonlinear programming problem (see Castonguay and Lasserre, 2019; Castonguay et al., 2020; Kling et al., 2016; Sanchirico and Springborn, 2011, for other applications of this technique). The dynamic controls to our problem—i.e. the drug and vaccine—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 or drug is licensed) to $t = T$ (assumed to be four months after the vaccine or drug is licensed) (Garg et al., 2010). 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. drug or vaccine capacity constraints) on state and control variables can be directly incorporated in the problem (Judd, 1998). 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 drugs and vaccines to fight COVID-19, this is likely due to the scarcity of the supply of vaccine and drugs 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 four control variables. The solution was found using TOMLAB (v. 8.4) (Holmström, 2001; Holmström et al., 2008) and the accompanying PROPT toolbox (Rutquist and Edvall, 2010). The approximate nonlinear programming problem is solved using general-purpose nonlinear optimization packages (e.g., KNITRO, SNOPT and NPSOL).

C Figures

C.1 Drugs

C.1.1 Compliance and Noncompliance to Travel Restrictions

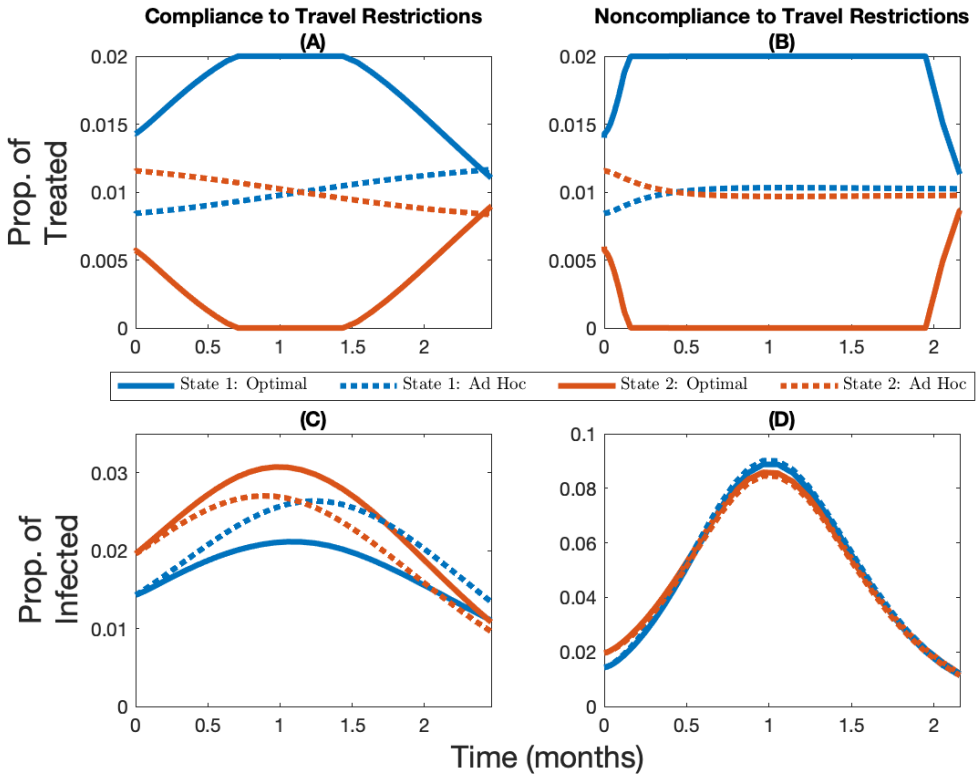

Figure A1: **Permanent immunity with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially 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 pharmaceutical intervention is an antiviral drug, the drug capacity constraint is 10%, and immunity is permanent.

C.1.2 Drug Capacity Constraints when Immunity is Permanent

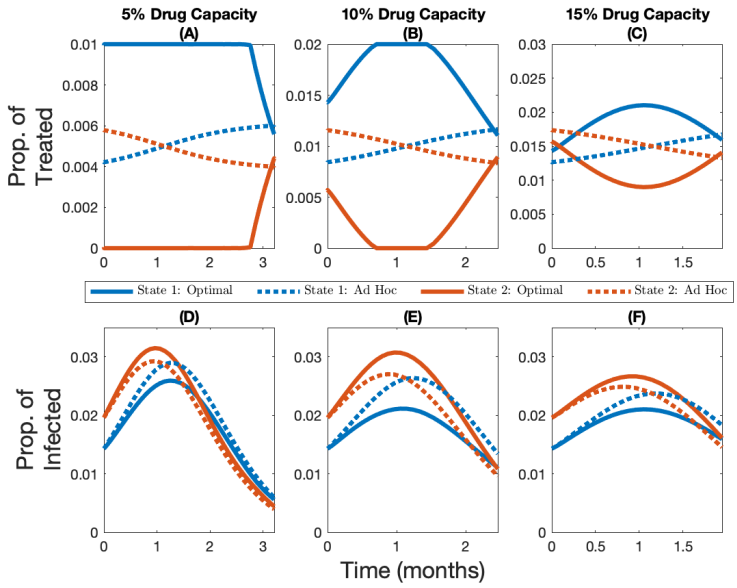

Figure A2: **Permanent immunity and compliance to travel restrictions with 5%, 10%, and 15% drug capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened state) 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 the pharmaceutical intervention is a antiviral drug, immunity is permanent, and there is compliance to travel restrictions.

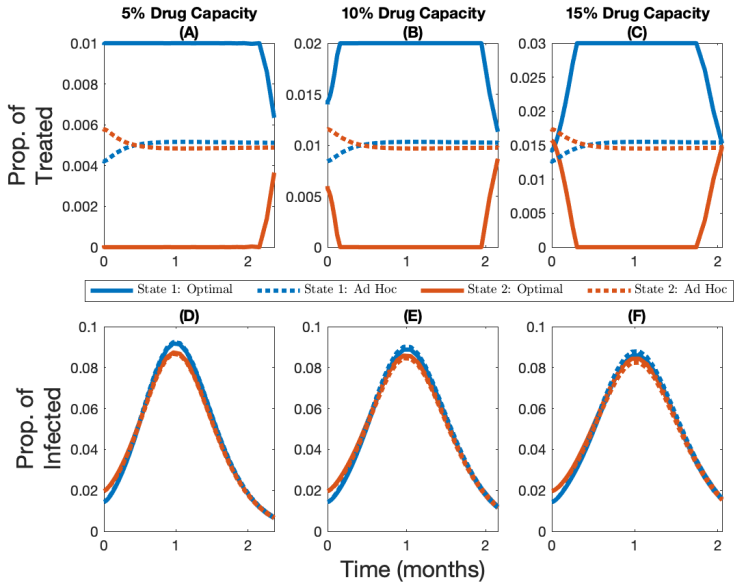

Figure A3: **Permanent immunity and noncompliance to travel restrictions with 5%, 10%, and 15% drug capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened 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 the pharmaceutical intervention is an antiviral drug, immunity is permanent, and there is no compliance to travel restrictions.

C.1.3 Drug Capacity Constraints when Immunity is Temporary

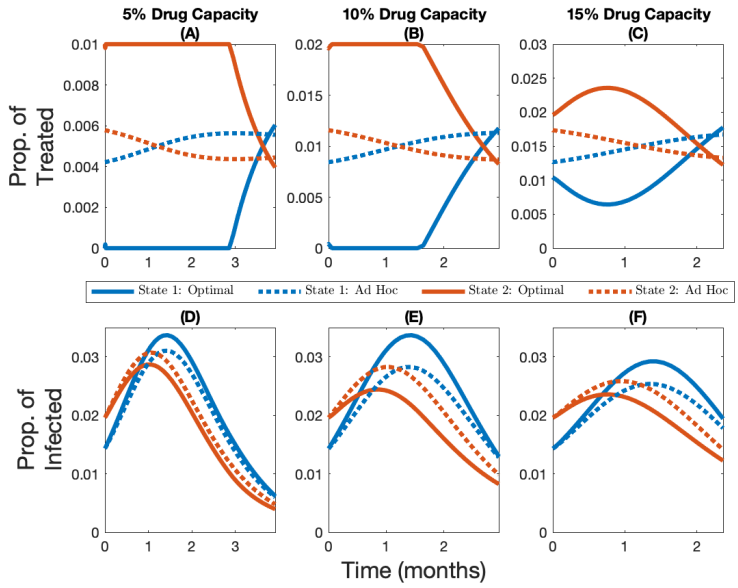

Figure A4: **Temporary immunity and compliance to travel restrictions with 5%, 10%, and 15% drug capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened state) 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 the pharmaceutical intervention is an antiviral drug, immunity lasts six months, and there is compliance to travel restrictions.

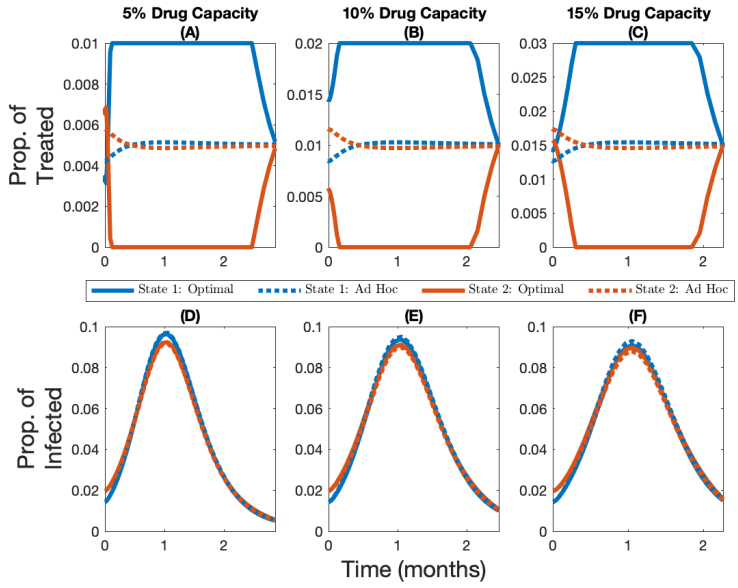

Figure A5: **Temporary immunity and noncompliance to travel restrictions with 5%, 10%, and 15% drug capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened 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 the pharmaceutical intervention is an antiviral drug, immunity lasts six months, and there is no compliance to travel restrictions.

C.1.4 Permanent vs Temporary Immunity

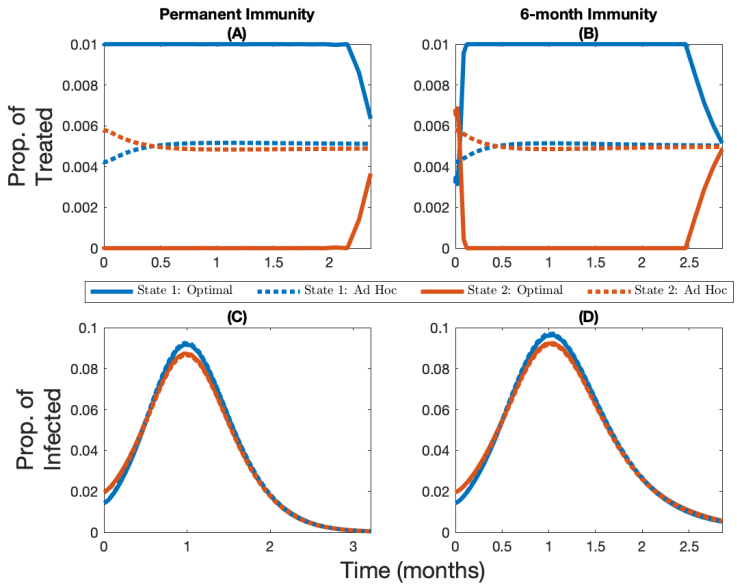

Figure A6: **Noncompliance to travel restrictions with permanent and temporary immunity.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, 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 pharmaceutical intervention is an antiviral drug, the drug capacity constraint is 5%, and there is no compliance to travel restrictions.

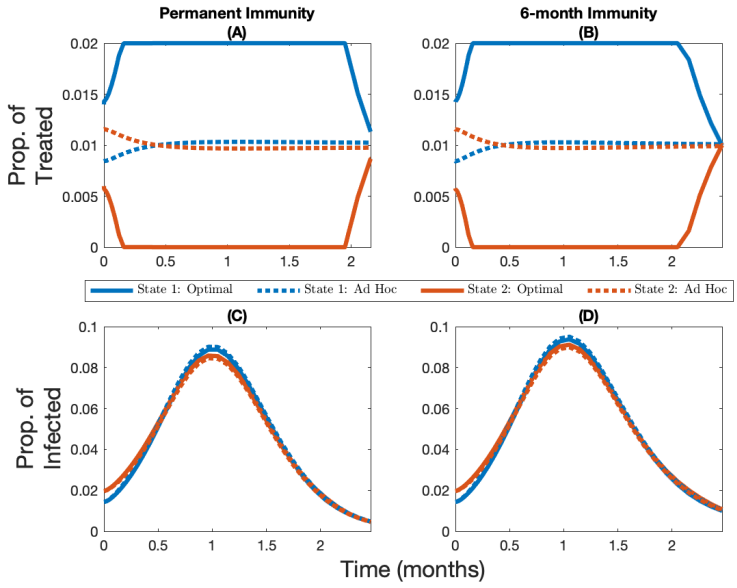

Figure A7: **Noncompliance to travel restrictions with permanent and temporary immunity.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, 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 pharmaceutical intervention is an antiviral drug, the drug capacity constraint is 10%, and there is no compliance to travel restrictions.

C.1.5 Cumulative Infection Levels

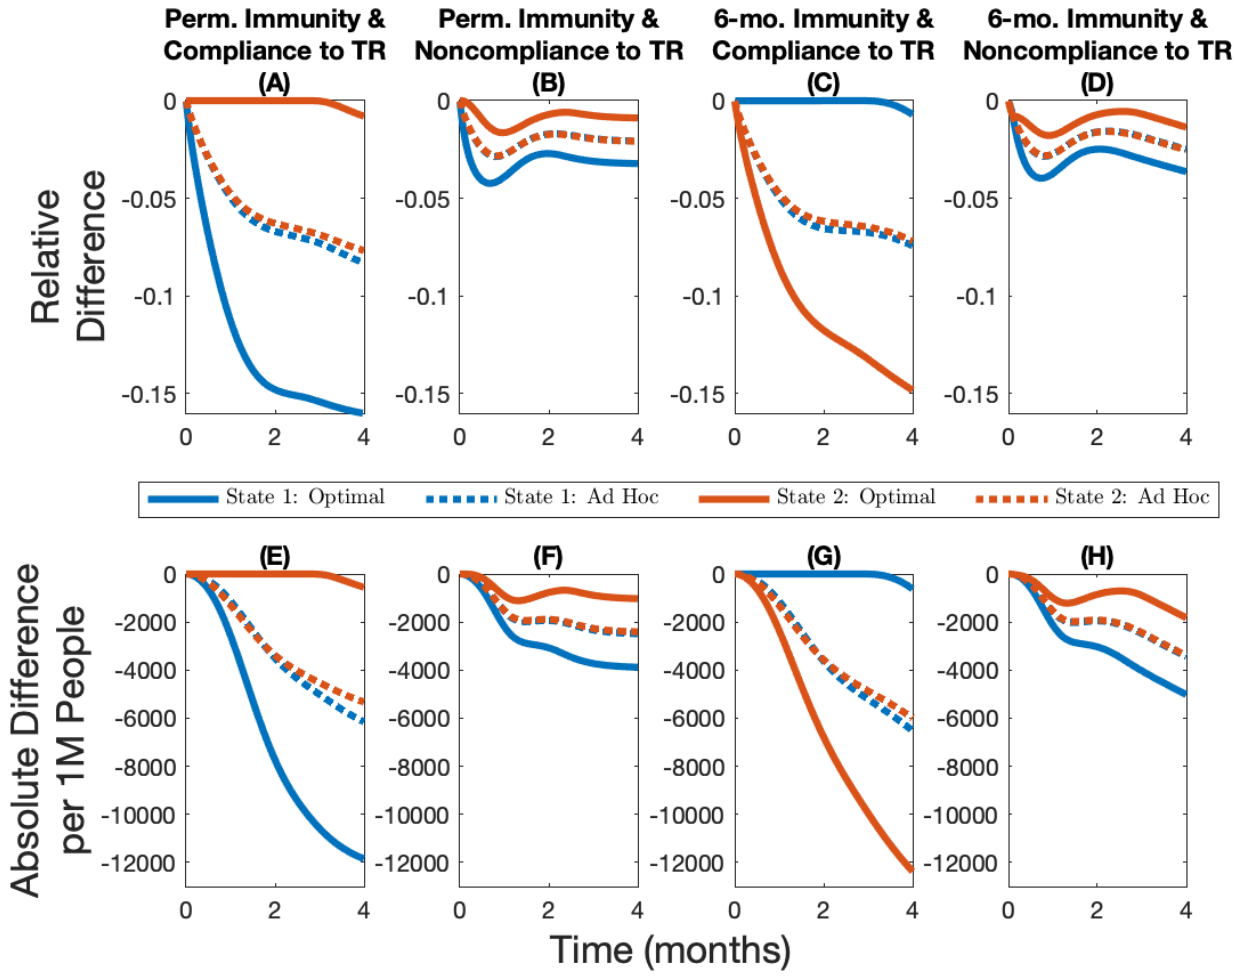

Figure A8: **Epidemiological outcomes under different scenarios with a low drug supply.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-drug case for different immunity–travel restrictions scenarios and for when drug capacity is 5%.

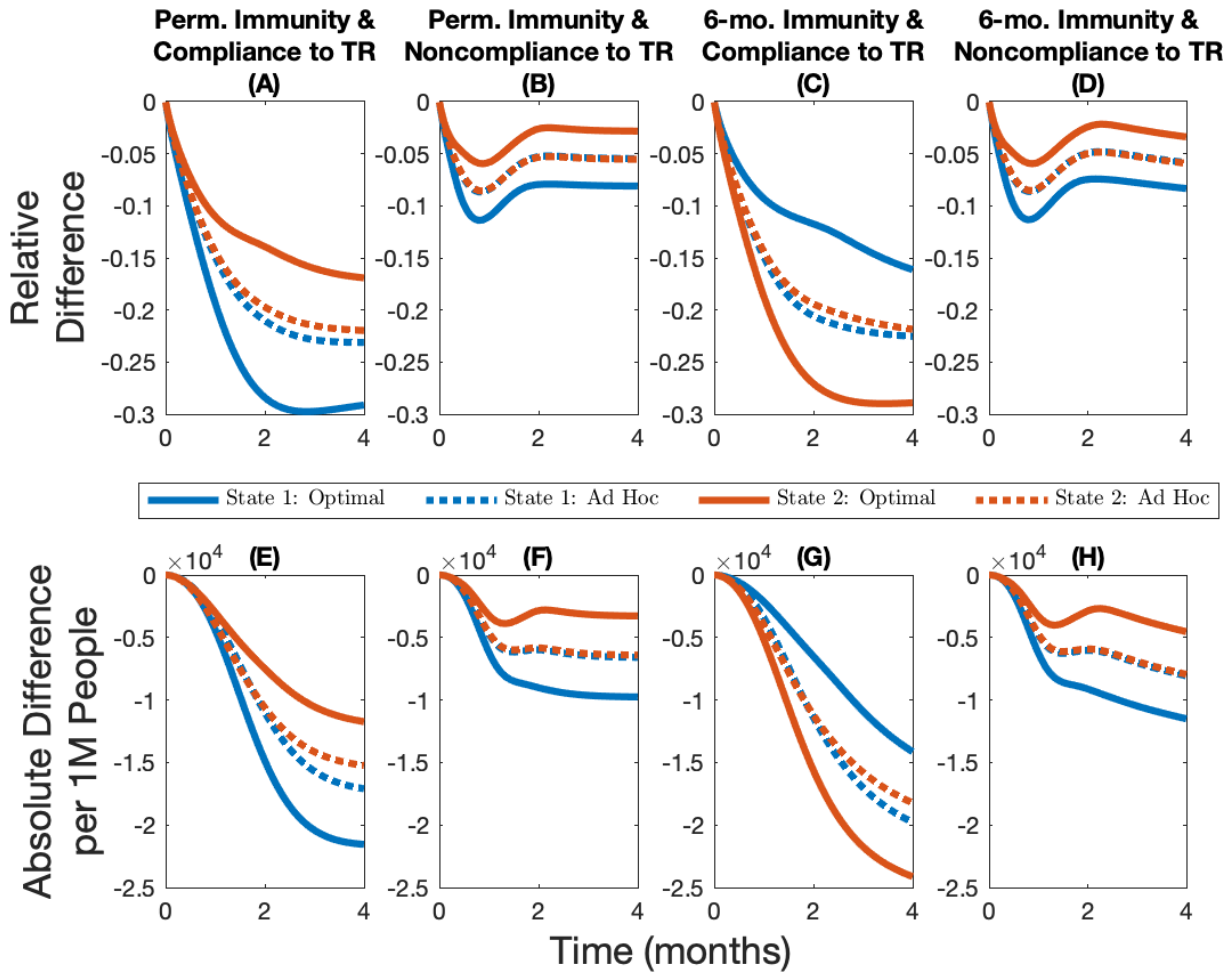

Figure A9: **Epidemiological outcomes under different scenarios with a high drug supply.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-drug case for different immunity–travel restrictions scenarios and for when drug capacity is 15%.

C.1.6 Robustness of Optimal Allocations

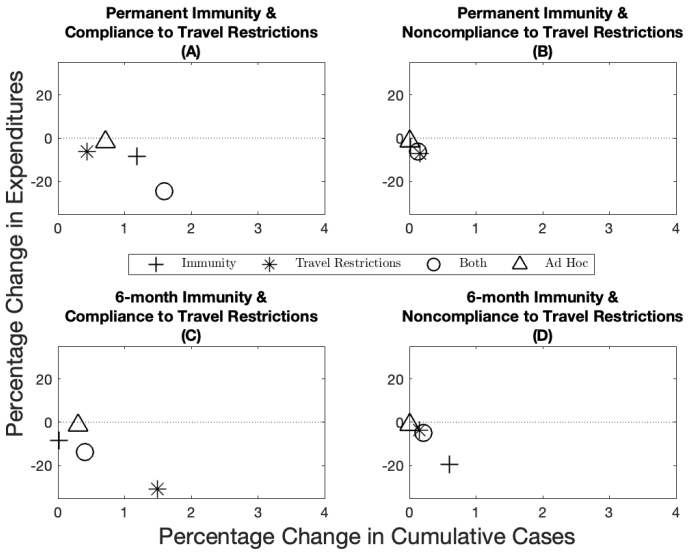

Figure A10: **Impact of incorrect assumption on epidemiological and economic outcomes under different scenarios with drugs.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 5%.

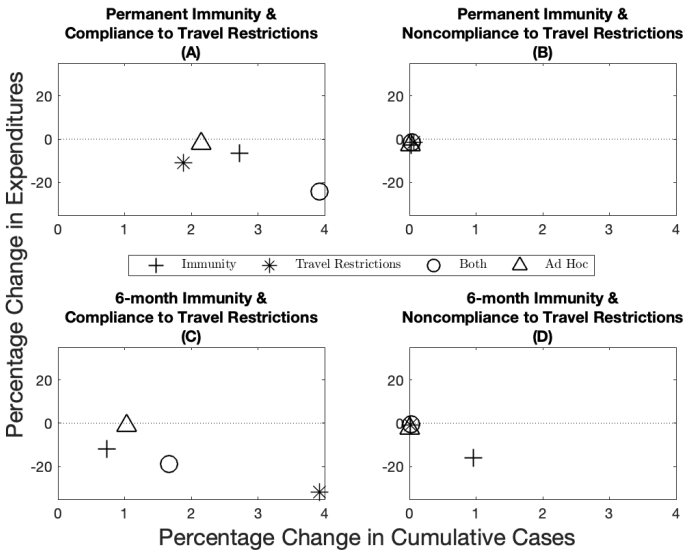

Figure A11: **Impact of incorrect assumption on epidemiological and economic outcomes under different scenarios with drugs.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 10%.

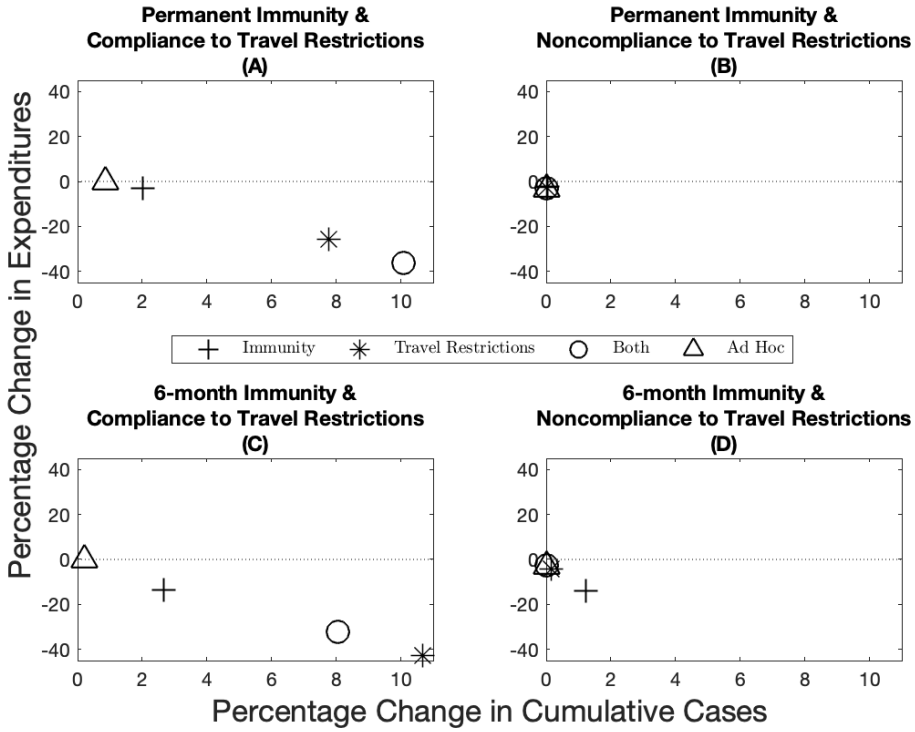

Figure A12: **Impact of incorrect assumption on epidemiological and economic outcomes under different scenarios with drugs.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 15%.

C.1.7 Sensitivity Analysis of Workability Cost

As mentioned above, imposing the *ad hoc* rules *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 or drug allocation over time with levels lower and higher than the base case parameter in the paper. We summarize these results by plotting the variance[19] of the optimal deviation in each time period from the *ad hoc* drug allocation (in blue; Figure A13 panels A, B, C, and D), and the difference in cumulative cases between the optimal and *ad hoc* allocation (in red; Figure A13 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 would expect that at lower values of the workability cost parameter we will find higher variance of the deviation. However, because the drug allocation gives rise to extreme cases where all the allotment is given to one jurisdiction for one period of time, the variance of the optimal deviation from the *ad hoc* cannot increase further. When the workability cost parameter reaches a magnitude in the neighborhood of the VSL (Figure A13 black line represents the VSL), then there is a discrete reduction in the variance and the optimal allocation converges to the *ad hoc* (in blue; Figure A13 panels A, B, C, and D) and any differences in cumulative cases disappear (in red; Figure A13 panels A, B, C, and D).

We also show how amount of funds allocated to the workability cost over time compare to expenditures on the total vaccine cost (Figure A13 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.

[19] The variance is calculated as Var(Optimal Drug − *Ad Hoc* Drug). Note that the variance of the optimal deviation from the *ad hoc* is identical in absolute terms across jurisdictions.

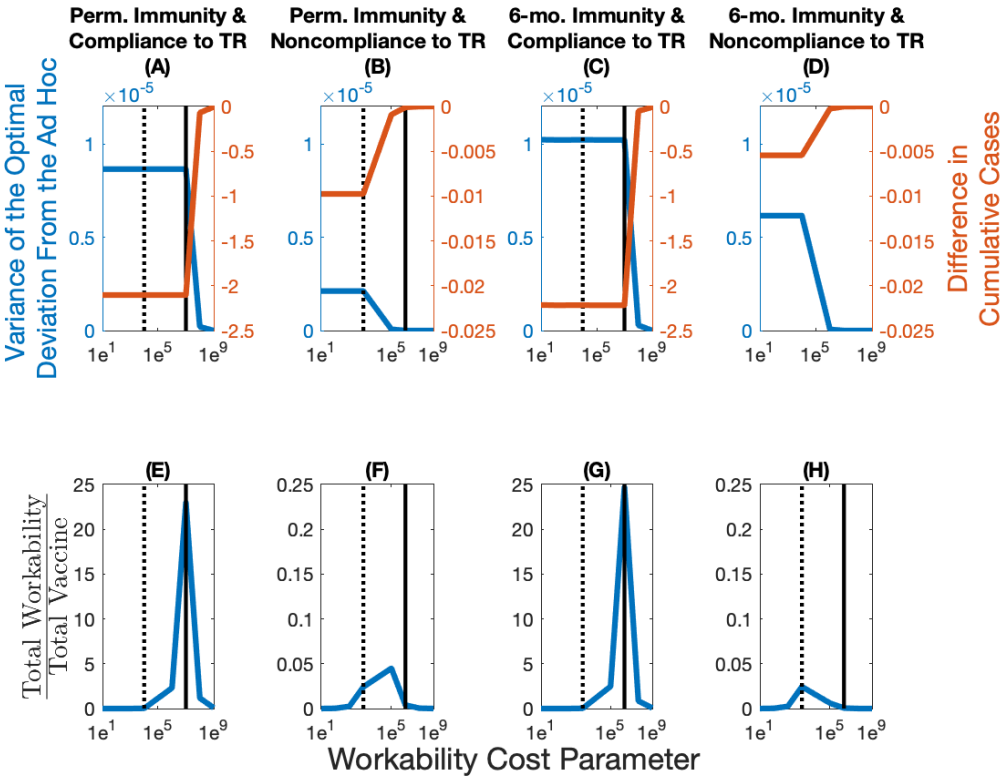

Figure A13: **Sensitivity of optimal allocations and epidemiological and economic outcomes when varying the workability cost parameter.** The variance of the optimal deviation in absolute terms (in blue; panels A, B, C, and D) represents an aggregate measure of the optimal deviation from the *ad hoc* 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.1.8 Sensitivity Analysis of Drug Effectiveness

The base case parameter for drug effectiveness we utilized in the paper is based on some low "clinically meaningfull" value (see Appendix A for more details). We investigate a range of what could be considered "clinically meaningful" and see how it impacts our results. When there is compliance to travel restrictions, we find that over a low range of effectiveness, the more a central planner would want to deviate from the *ad hoc* allocation. For a high range of effectiveness, however, the deviation starts decreasing because the drug becomes so effective that the physical constraints are slack (i.e. there are fewer infected individuals that there are available drugs) and the optimal drug allocation converges towards the *ad hoc* (in blue; Figure A14 panels A and C). When there is noncompliance to travel restrictions, we find that for the entire range of drug effectiveness considered, the more effective the drug is, the more a central planner would want to deviate from the *ad hoc* allocation (in blue; Figure A14 panels B, and D). Regardless of the situation, a higher drug effectiveness implies a larger difference in terms of the reduction in cumulative cases (in red; Figure A14 panels A, B, C, and D). When there is compliance to travel restrictions, both the total workability cost and total vaccine cost decrease at high values of drug effectiveness; the total workablity cost decreases relatively faster than the total vaccine costs which together implies a lower ratio with higher levels of drug effectiveness (Figure A14 panels E, F, G, and H).

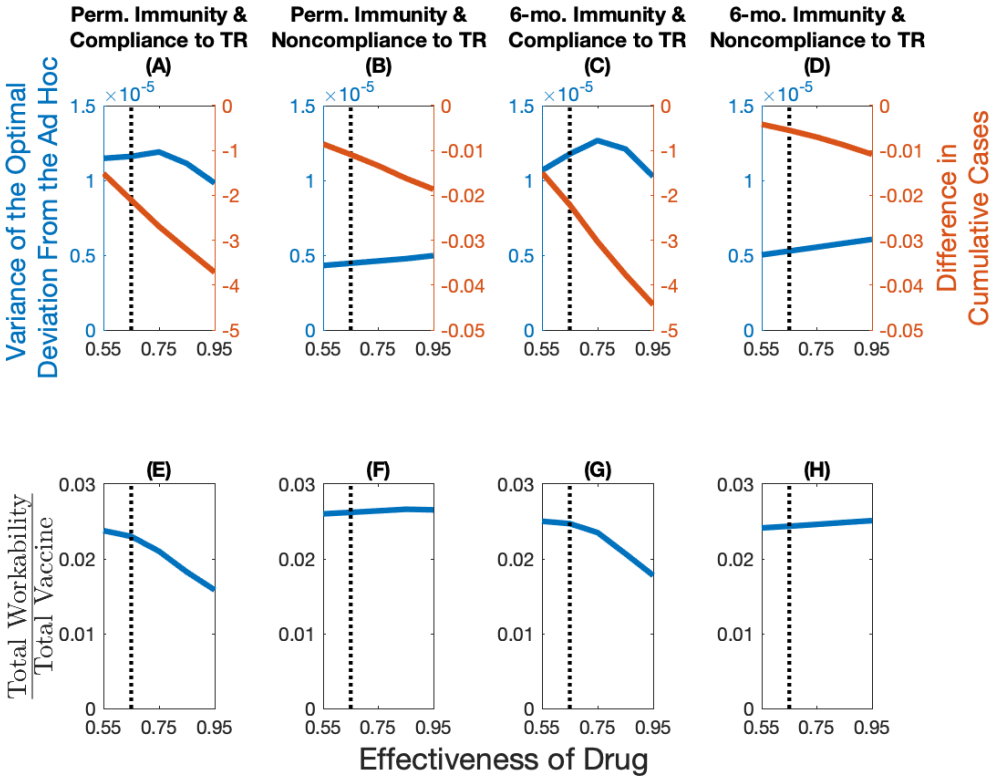

Figure A14: **Sensitivity of optimal allocations, and of epidemiological and economic outcomes when varying the effectiveness of the antiviral drug.** The variance of the optimal deviation in absolute terms (in blue; panels A, B, C, and D) represents an aggregate measure of the optimal deviation from 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 drug 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).

C.2 Vaccines

C.2.1 Compliance and Noncompliance to the Travel Restrictions

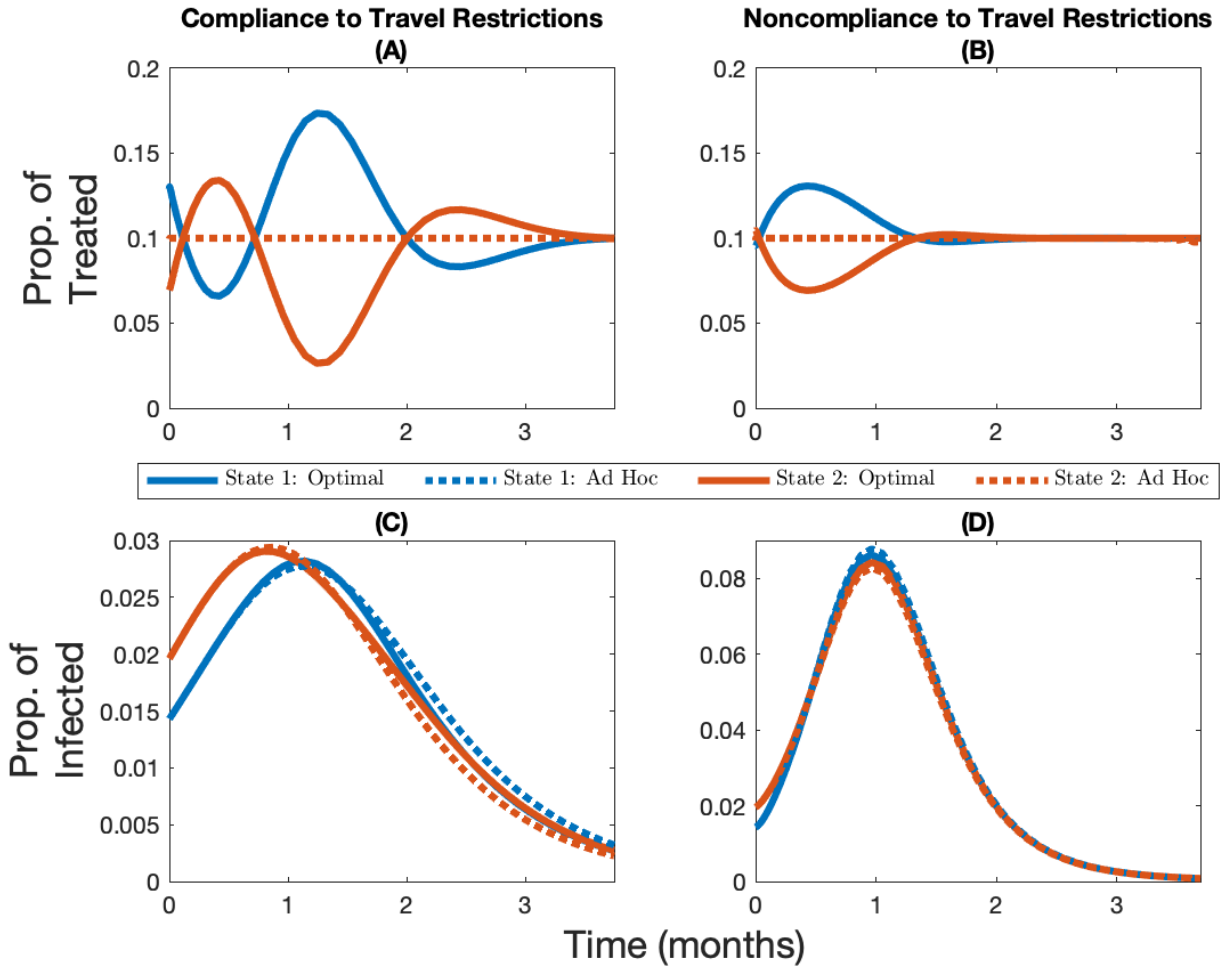

Figure A15: **Permanent immunity with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the 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 pharmaceutical intervention is a vaccine, the vaccine capacity constraint is at 5% of the unconstrained case, and immunity lasts six months.

C.2.2 Vaccine Capacity Constraints when Immunity is Permanent

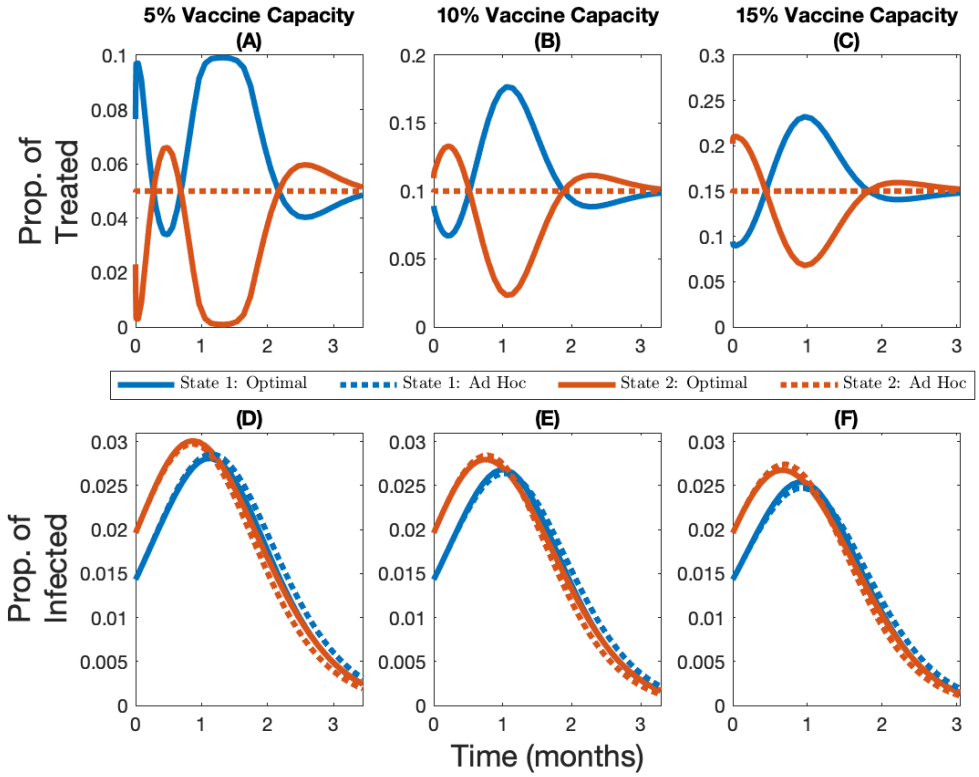

Figure A16: **Permanent immunity and compliance to travel restrictions with 5%, 10%, and 15% vaccine capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened 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 the pharmaceutical intervention is a vaccine, immunity is permanent, and there is compliance to travel restrictions.

C.2.3 Vaccine Capacity Constraints when Immunity is Temporary

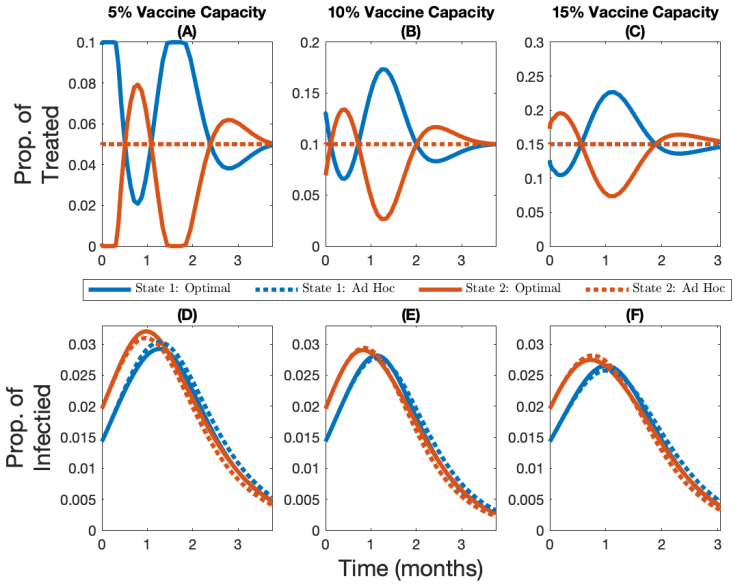

Figure A17: **Temporary immunity and compliance to travel restrictions with 5%, 10%, and 15% vaccine capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened 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 the pharmaceutical intervention is an antiviral drug, immunity lasts six months, and there is compliance to travel restrictions.

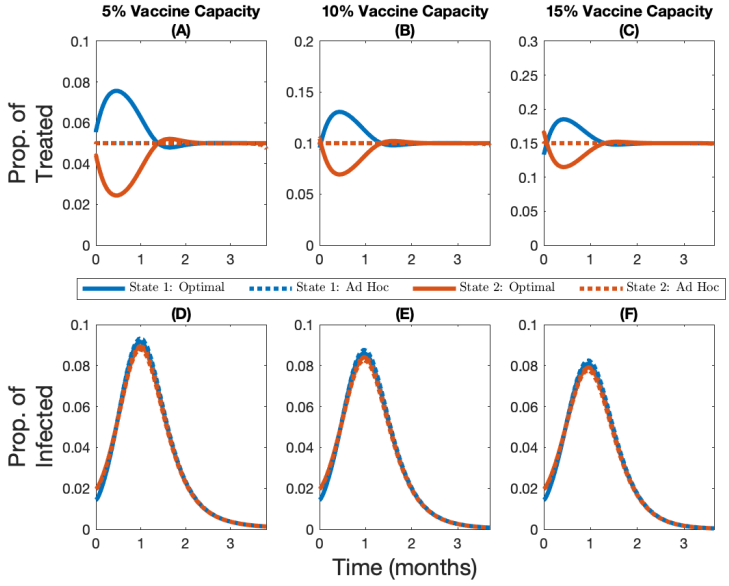

Figure A18: **Temporary immunity and noncompliance to travel restrictions with 5%, 10%, and 15% vaccine capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened 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 the pharmaceutical intervention is an antiviral drug, immunity lasts six months, and there is no compliance to travel restrictions.

C.2.4 Permanent vs Temporary Immunity

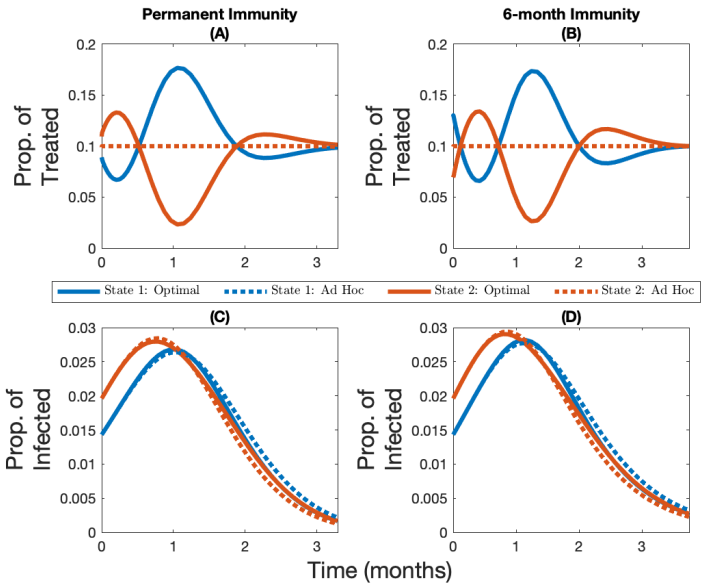

Figure A19: **Compliance to travel restrictions with permanent and temporary immunity.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the 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 pharmaceutical intervention is a vaccine, the vaccine capacity constraint is 10%, and there is compliance to travel restrictions.

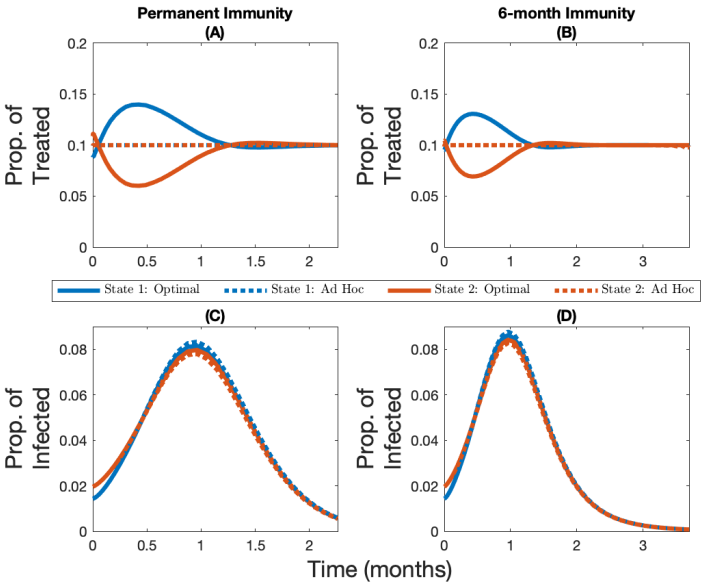

Figure A20: **Noncompliance to travel restrictions with permanent and temporary immunity.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, 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 pharmaceutical intervention is a vaccine, the vaccine capacity constraint is 10%, and there is no compliance to travel restrictions.

C.2.5 Cumulative Infection Levels

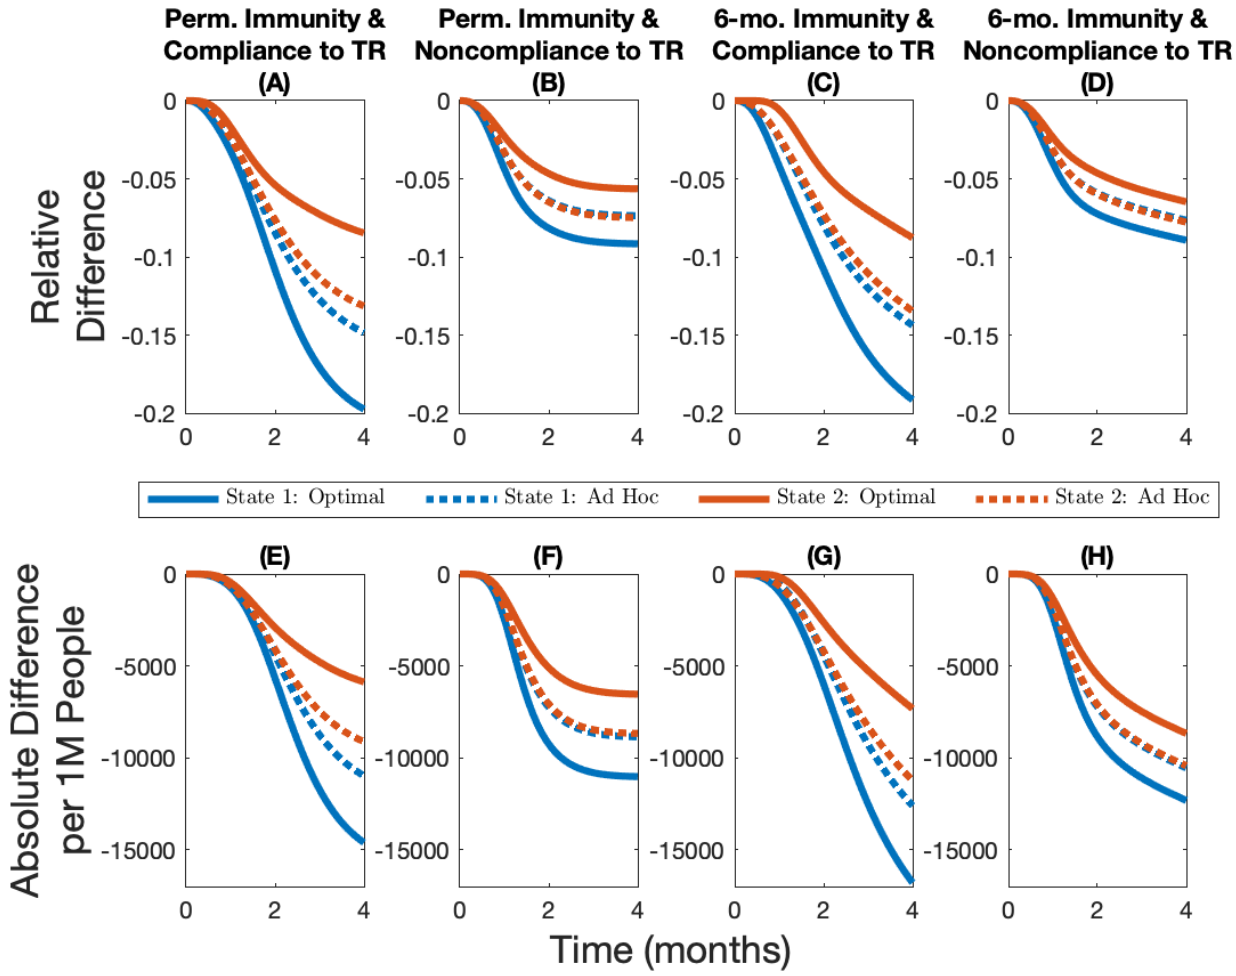

Figure A21: **Epidemiological outcomes under different scenarios with a low vaccine supply.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-vaccine case for different immunity–travel restrictions scenarios and for when vaccine capacity is 5%.

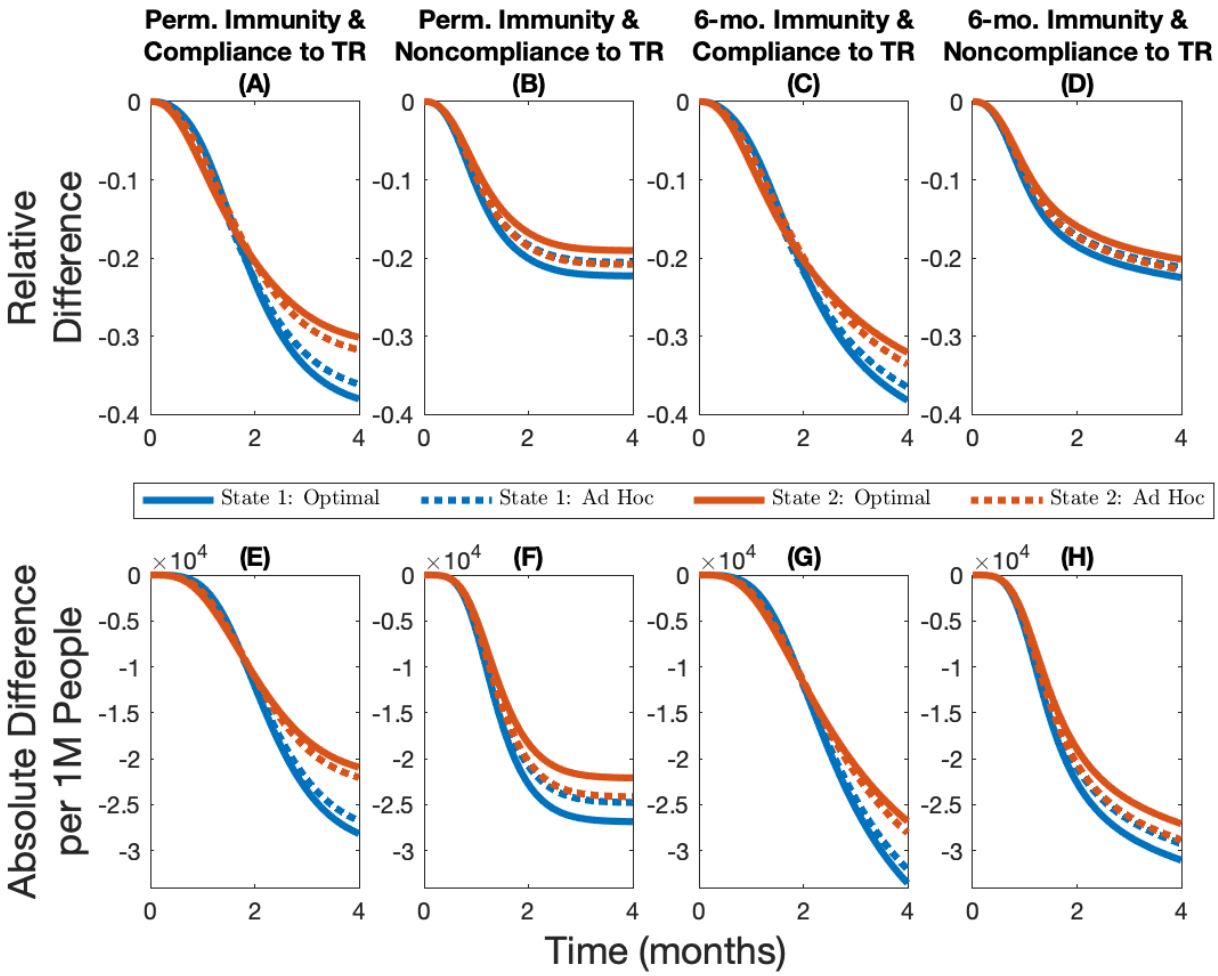

Figure A22: **Epidemiological outcomes under different scenarios with a high vaccine supply.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-vaccine case for different immunity–travel restrictions scenarios and for when vaccine capacity is 15%.

C.2.6 Robustness of Optimal Allocations

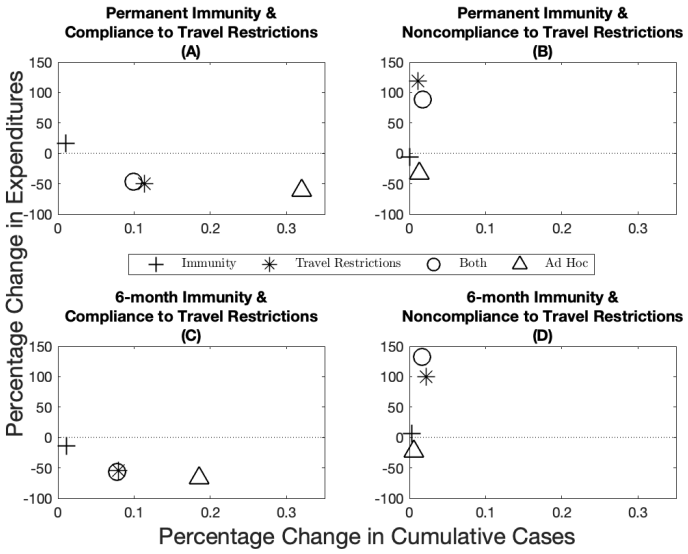

Figure A23: **Impact of incorrect assumption on epidemiological and economic outcomes under different scenarios with vaccines.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 5%.

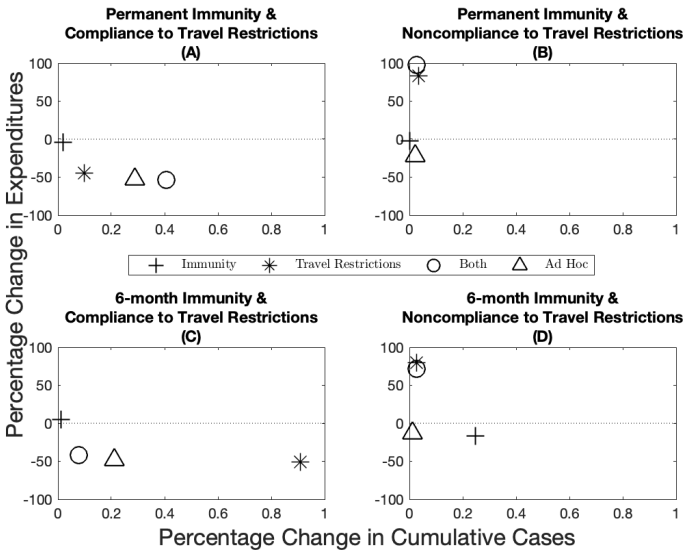

Figure A24: **Impact of incorrect assumption on epidemiological and economic outcomes under different scenarios with vaccines.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 15%.

C.2.7 Sensitivity Analysis of Workability Cost

As mentioned above, imposing the *ad hoc* rules *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 or drug allocation over time with levels lower and higher than the base case parameter in the paper. We summarize these results by plotting the variance[20] of the optimal deviation in each time period from the *ad hoc* vaccine allocation (in blue; Figure A25 panels A, B, C, and D), and the difference in cumulative cases between the optimal and *ad hoc* allocation (in red; Figure A25 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 A25 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 A25 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.

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

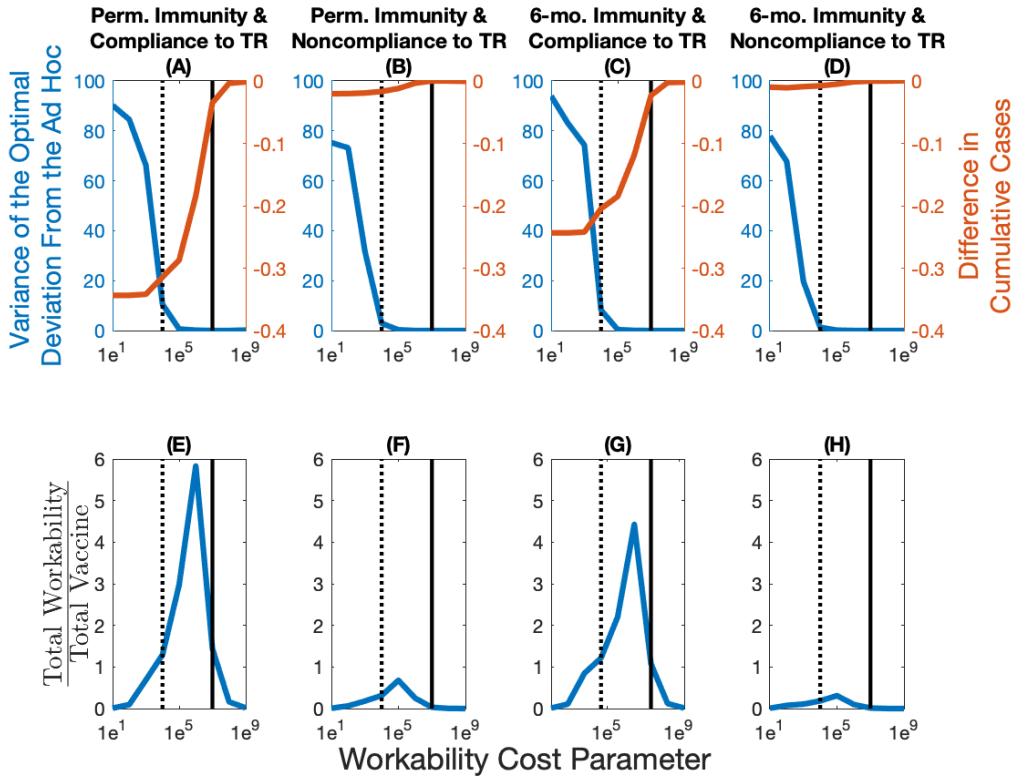

Figure A25: **Sensitivity of optimal allocations and epidemiological and economic outcomes when varying the workability cost parameter.** The variance of the optimal deviation in percentage (in blue; panels A, B, C, and D) represents an aggregate measure of the optimal deviation from the *ad hoc* allocation. 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.2.8 Sensitivity Analysis of Vaccine Effectiveness

The base case parameter for vaccine effectiveness we utilized in the paper is based on estimates of the influenza vaccine (Ohmit et al., 2014, 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 A26 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 A26 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 A26 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.

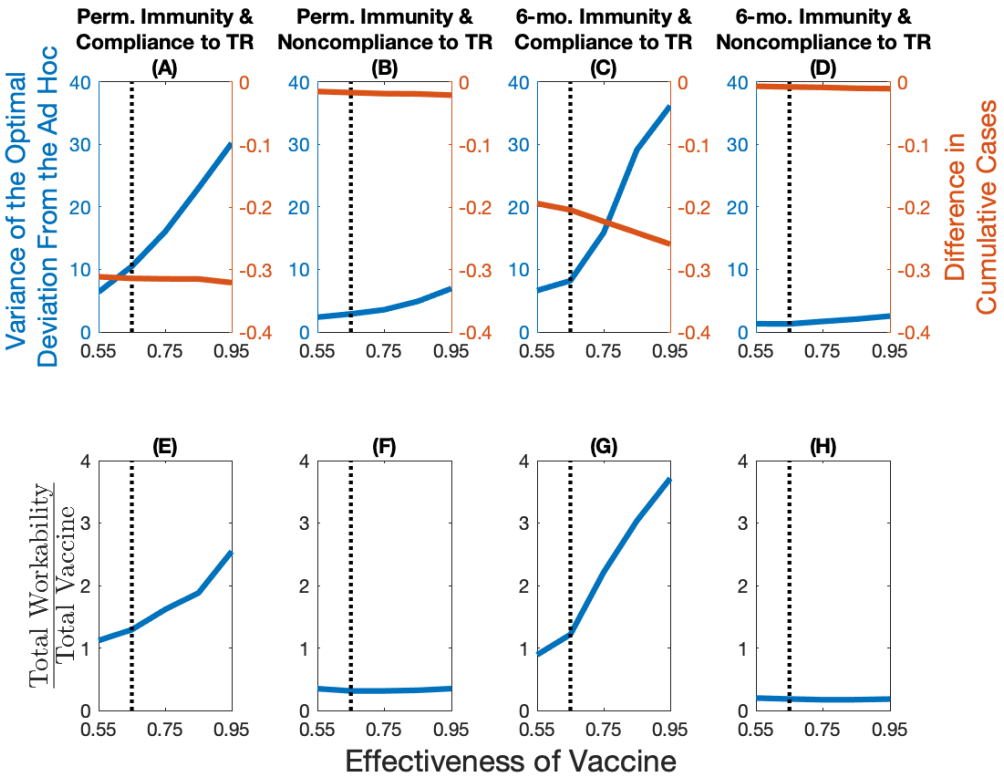

Figure A26: **Sensitivity of optimal allocations, and of epidemiological and economic outcomes when varying the effectiveness of the vaccine.** The variance of the optimal deviation in percentage (in blue; panels A, B, C, and D) represents an aggregate measure of the optimal deviation from the *ad hoc* 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

Abdollahi, E., Champredon, D., Langley, J. M., Galvani, A. P., and Moghadas, S. M. (2020). Temporal estimates of case-fatality rate for COVID-19 outbreaks in Canada and the United States. *CMAJ*.

Barrett, S. and Hoel, M. (2007). Optimal disease eradication. *Environment and Development Economics*, pages 627–652.

Buckner, J. H., Chowell, G., and Springborn, M. R. (2020). Optimal dynamic prioritization of scarce COVID-19 vaccines. *medRxiv*.

Castonguay, F. M. and Lasserre, P. (2019). L'exploitation de ressources naturelles non renouvelables en asymétrie d'information. *L'Actualité économique*, 95(2-3).

Castonguay, F. M., Sokolow, S. H., De Leo, G. A., and Sanchirico, J. N. (2020). Cost-effectiveness of combining drug and environmental treatments for environmentally transmitted diseases. *Proceedings of the Royal Society B*, 287(1933):20200966.

Davies, N. G., Klepac, P., Liu, Y., et al. (2020). Age-dependent effects in the transmission and control of COVID-19 epidemics. *Nature Medicine*, 26(8):1205–1211.

Diekmann, O., Heesterbeek, J., and Roberts, M. G. (2010). The construction of next-generation matrices for compartmental epidemic models. *Journal of the Royal Society Interface*, 7(47):873–885.

Diekmann, O., Heesterbeek, J. A. P., and Metz, J. A. (1990). On the definition and the computation of the basic reproduction ratio R0 in models for infectious diseases in heterogeneous populations. *Journal of Mathematical Biology*, 28(4):365–382.

Edridge, A. W., Kaczorowska, J. M., Hoste, A. C., Bakker, M., Klein, M., Jebbink, M. F., Matser, A., Kinsella, C., Rueda, P., Prins, M., et al. (2020). Coronavirus protective immunity is short-lasting. *MedRxiv*.

Environmental Protection Agency (2020). *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).

Garchitorena, A., Sokolow, S., Roche, B., Ngonghala, C., Jocque, M., Lund, A., Barry, M., Mordecai, E., Daily, G., Jones, J., et al. (2017). 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*, 372(1722):20160128.

Garg, D., Patterson, M., Hager, W. W., Rao, A. V., Benson, D. A., and Huntington, G. T. (2010). A unified framework for the numerical solution of optimal control problems using pseudospectral methods. *Automatica*, 46(11):1843–1851.

Goldman, S. M. and Lightwood, J. (2002). Cost optimization in the sis model of infectious disease with treatment. *The BE Journal of Economic Analysis & Policy*, 2(1).

Holmström, K. (2001). Practical optimization with the tomlab environment in matlab. In *Proceedings of the 42nd SIMS Conference, Sep 15, 2001*, pages 89–108. Citeseer. Porsgrunn, Norway: Telemark University College.

Holmström, K., Göran, A. O., and Edvall, M. M. (2008). Users guide for tomlab/snopt. *Mälardalen University, Department of Mathematics and Physics, Västerås, Sweden*.

John, J., Koerber, F., and Schad, M. (2019). Differential discounting in the economic evaluation of healthcare programs. *Cost Effectiveness and Resource Allocation*, 17(1):29.

Judd, K. L. (1998). *Numerical methods in economics*. MIT press. Cambridge, Massachusetts.

Kling, D. M., Sanchirico, J. N., and Wilen, J. E. (2016). Bioeconomics of managed relocation. *Journal of the Association of Environmental and Resource Economists*, 3(4):1023–1059.

Li, Q., Guan, X., Wu, P., Wang, X., Zhou, L., Tong, Y., Ren, R., Leung, K. S., Lau, E. H., Wong, J. Y., et al. (2020). Early transmission dynamics in wuhan, china, of novel coronavirus–infected pneumonia. *New England Journal of Medicine*.

Nurchis, M. C., Pascucci, D., Sapienza, M., Villani, L., D'Ambrosio, F., Castrini, F., Specchia, M. L., Laurenti, P., and Damiani, G. (2020). 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*, 17(12):4233.

Ohmit, S. E., Thompson, M. G., Petrie, J. G., Thaker, S. N., Jackson, M. L., Belongia, E. A., Zimmerman, R. K., Gaglani, M., Lamerato, L., Spencer, S. M., et al. (2014). Influenza vaccine effectiveness in the 2011–2012 season: protection against each circulating virus and the effect of prior vaccination on estimates. *Clinical infectious diseases*, 58(3):319–327.

Rowthorn, R. E., Laxminarayan, R., and Gilligan, C. A. (2009). Optimal control of epidemics in metapopulations. *Journal of the Royal Society Interface*, 6(41):1135–1144.

Rutquist, P. E. and Edvall, M. M. (2010). Propt-Matlab optimal control software. *Tomlab Optimization Inc*, 260(1).

Ryan, D., Toews, C., Sanchirico, J. N., and Armsworth, P. R. (2017). Implications of policy adjustment costs for fisheries management. *Natural Resource Modeling*, 30(1):74–90.

Sanchirico, J. N. and Springborn, M. (2011). How to get there from here: ecological and economic dynamics of ecosystem service provision. *Environmental and Resource Economics*, 48(2):243–267.

Schlosser, F., Maier, B. F., Jack, O., Hinrichs, D., Zachariae, A., and Brockmann, D. (2020). Covid-19 lockdown induces disease-mitigating structural changes in mobility networks. *Proceedings of the National Academy of Sciences*.

Tian, H., Liu, Y., Li, Y., Wu, C.-H., Chen, B., Kraemer, M. U., Li, B., Cai, J., Xu, B., Yang, Q., et al. (2020). An investigation of transmission control measures during the first 50 days of the COVID-19 epidemic in china. *Science*, 368(6491):638–642.

U.S. Department of Health and Human Services, Food and Drug Administration (FDA), and Center for Biologics Evaluation and Research (CBER) (2020). Development and licensure of vaccines to prevent COVID-19: guidance for industry. june, 2020. Technical report, Food and Drug Administration, 5630 Fishers Lane, Rm. 1061, Rockville, MD 20852.

U.S. Department of Health and Human Services, Food and Drug Administration (FDA), Center for Biologics Evaluation and Research (CBER), and Center for Drug Evaluation and Research (CDER) (2019). Demonstrating substantial evidence of effectiveness for human drug and biological products: Guidance for industry. Technical report, Food and Drug Administration, 5630 Fishers Lane, Rm. 1061, Rockville, MD 20852.