Economic and Health Consequences of Inefficient Treatment of Communicable Human Diseases: Applications to Environmentally Transmitted Diseases and Vaccine-Preventable Diseases

By

FRANCOIS CASTONGUAY

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Agricultural and Resource Economics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

James N. Sanchirico, Chair

James E. Wilen

Michael R. Springborn

Committee in Charge

2021

Acknowledgments

I want to start this dissertation by thanking everyone who supported me over the years, and without whom this journey would have been completely different and oh-so much more difficult. Today, I am completing an 11-year chapter of my life and I am about to start a new one with the Health Economics and Modeling Unit (HEMU) of the National Center for Emerging and Zoonotic Infectious Diseases (NCEZID) at the Centers for Disease Control and Prevention (CDC), amid the Covid-19 pandemic.

I want to thank my family and friends back in Quebec and Connecticut. Even from afar, your unconditional support has been very important to me every step of the way. I also want to acknowledge the Department of Agricultural and Resource Economics at UC Davis without whom none of this would have been possible; thank you for offering me this opportunity to surpass myself. The colleagues—nay, friends!—I have had the chance to meet and interact with over these last five years have helped me tremendously in overcoming my fears and doubts, and have been an indispensable source of fun and laughs, which gave me the strength to push forward.

Another primordial piece of this 11-year journey has been all the help I have gotten from my professors. More particularly, from my master's at UQAM, I want to thank the late Professor Pierre Lasserre. Pierre, ton départ soudain me laisse à ce jour attristé et sans mots ; sache qu'indirectement tu ne cesseras jamais de m'aider professionnellement. From my time at UC Davis, I want to thank the professors James N. Sanchirico, James E. Wilen, and Michael R. Springborn. Jim W., I feel privileged to have had the opportunity to learn and interact with you; your wisdom and advice have been invaluable. Mike, our interactions have always been fruitful and helpful for my research—every time, you bring a unique and fresh perspective to my work.

Jim S., your help and support over the last four years has been crucial and determinant. I will never find the right words to properly thank you for everything you have done for me, but I still wanted you to know the enormous impact you had on me and my professional development. Without you, I would not be writing these words today.

Abstract

Communicable diseases spread from human to human through viruses, bacteria, parasites, and fungi via direct contact with an infected individual, the individual's discharge, or by indirect means (i.e. vectors). Treatment of communicable diseases is simultaneously a private good by reducing disability in treated individuals and a public good by reducing the likelihood of disease transmission between infected and susceptible individuals. Due to their inherent nature, optimal treatment of communicable diseases is difficult, if not impossible, to achieve without intervention from a social planner. In the three chapters of this dissertation, I identify the sources of potential inefficiencies that may occur in the treatment of environmentally transmitted diseases without vaccines and an emerging disease that has a newly developed vaccine. In the first chapter, I highlight how treatment inefficiencies may arise when there are several ways of treating the disease. I show that when recommendations of the different treatments are determined independently, there is excessive usage of public funds and overutilization of treatments. In the second chapter, I show the benefits of targeting public health interventions to specific geographical areas—which may differ for various reasons such as the timing of the outbreak, the age structure of the population, or the number of essential workers—when resources are scarce, rather than using some *ad hoc* rule based on, for instance, relative population size. I also investigate how robust the optimal allocation is when allocation decisions must be made before uncertainty is resolved (e.g. before the duration of immunity is known). In the third chapter, I examine potential unintended consequences that may arise when the main driver of some public health intervention is the economic benefits indirectly derived by it; this shows the importance of considering the impact of human behavior when designing public health policies. As a whole, this dissertation uses various examples to show why treatment of communicable diseases requires intervention from a social planner and how inefficient disease management can lead to dire economic and health consequences. It also offers different context-specific avenues to mitigate the impact of communicable diseases.

1 Chapter 1: "Cost-Effectiveness of Combining Drug and Environmental Treatments for Environmentally Transmitted Diseases"

Abstract

Control of neglected tropical diseases (NTDs) via mass drug administration (MDA) increased considerably over the past decade but strategies focused exclusively on human treatment show limited efficacy. This paper investigated trade-offs between drug and environmental treatments in the fight against NTDs by using schistosomiasis as a case study. We use optimal control techniques where the planner's objective is to treat the disease over a time horizon at the lowest possible total cost, where the total costs include treatment, transportation, and damages (reduction in human health). We show that combining environmental treatments and drug treatments reduces the dependency on MDAs and that this reduction increases when the planners take a longer run perspective on the fight to reduce NTDs. Our results suggest that NTDs with environmental reservoirs require moving away from a reliance solely on MDA to integrated treatment involving investment in both drug and environmental controls.

Key words: economic epidemiology, applied numerical methods, efficient treatment of diseases, misallocation of resources.

1.1 Introduction

Neglected tropical diseases (NTDs) affect approximately one in six people, mainly in the poorest rural and remote areas, urban slums, and conflict zones. The loss of disability-adjusted life years due to NTDs (48 million) is as high as tuberculosis (49 million), and more than half of malaria (83 million) and HIV/AIDS (82 million) [1]. NTDs also increase the risk of coinfection; they are responsible for one-half and one-third of sub-Saharan Africa's malaria and HIV/AIDS disease burden, respectively [2]. Despite major donations

from pharmaceutical companies, private foundations (e.g., Bill and Melinda Gates) and foreign governments (e.g., UK and US), the World Health Organization (WHO) estimated that an additional 2 billion US dollars was needed to administer preventive chemotherapy to all individuals who were at risk of contracting an NTD between 2012 and 2015 [3].

A subset of NTDs are environmentally transmitted diseases (ETDs) where pathogens rely partially, or entirely, on non-human hosts, reservoirs or vectors. ETDs with focal transmission exhibit a direct link between the infection rate in the hosts and the level of the pathogen in an environmental reservoir [4]. The link can create an inherent cycle between the population's infection rate and the environmental degradation of contaminated reservoirs.

Schistosomiasis, which is a focus of this paper, is an example of an ETD with focal transmission [4]. The global disease burden of schistosomiasis has remained relatively stable despite the development almost a half century ago of an anthelmintic drug, praziquantel, that promised widespread control. In the last two decades, over 1.4 billion US dollars was spent on a mass drug administration (MDA) treatment protocol for schistosomiasis [5].

Completely eliminating schistosomiasis' pathogen transmission seems difficult, if not impossible, to achieve. The WHO's guidelines [6] recommend to target school-age children, given the facility to deliver treatment in schools (see, e.g., [7; 8; 9; 10]), with community wide treatment (i.e., including adults) being recommended in high prevalence communities (see, e.g., [11; 12]). Elimination is very challenging, because targeting children reduces, but does not eliminate the shedding of pathogens into the reservoir (see Figure 1 for different treatment alternatives and how they interrupt the life-cycle of the pathogen). Even if the whole community could be treated, shedding of pathogen into the environment remains due to noncompliance to drug treatment [13; 14; 15] and limited effectiveness of drug controls [16; 17]. Furthermore, the fact that individuals treated via MDA often have no other alternative but to return to parasite-contaminated waters [4] means that reinfection is likely to occur.

Although most treatment of NTDs consists of implementing MDAs [6], there is increasing evidence for focally-transmitted ETDs that water quality, sanitation and hygiene (WASH) measures and environmental treatments (EnvTre) can have significant positive impacts on health outcomes by reducing contamination (e.g., via sanitation measures) and exposure to pathogens (e.g., by providing safe water) (see Andres et al. [18] for a meta-analysis of WASH impact evaluations). Environmental treatments act in a similar manner to WASH by reducing the transmission pathways between the disease reservoir and human contact but focus more directly on reducing pathogen abundance in the reservoir or preventing transmission from the reservoir to humans (Figure 1). For example, an environmental treatment, EnvTre for short, can reduce reservoir, vector or intermediate host populations (e.g., chemical molluscicides or insecticides) or reduce the pathogens directly via targeting their free-living stages in water or soil (e.g., chlorination).

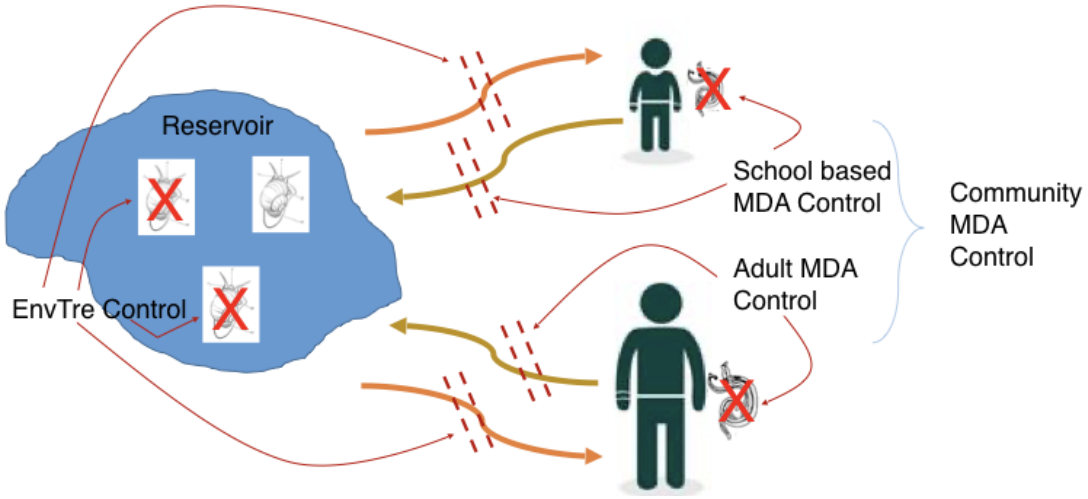

Figure 1: Modeled treatment options and transmission pathways for schistosomiasis.

The combination of multiple types of treatment for schistosomiasis and other ETDs can potentially reduce the overall cost associated with treating the disease and reduce the disease burden [5; 12; 19; 20]. For example, Lo et al. [12] demonstrate a cost-effective combination of controls that reduce the prevalence of the shistosomiasis pathogen in the environment with school-based MDA treatment. Most of the current literature considers combinations of controls using simulation (scenario) analysis under the assumption of a fixed level of MDA treatment occurring on a pre-determined set interval (e.g., every

other year) and a fixed level of a WASH or environmental treatment on a similarly fixed interval (not necessarily the same as the MDA treatment interval). In these analyses, understanding the when, where, and how much to combine to achieve the most cost-effective combination is challenging due to potential direct and indirect effects of one type of treatment on another and all of the possible combinations of multiple treatments available over time.

Our paper makes a number of important contributions to the literature on treatment for ETDs and specifically shistosomiasis control. First, we consider multiple combinations of treatments that include MDA and environmental controls in an optimal control framework that solves for the optimal cost-effective solution (for other applications of optimal control to schistosomiasis, see [21; 22; 23; 24]). We use optimal control to examine the use of both MDA and environmental treatments, and to understand under what conditions both approaches should be used in combination or in series. Our methodological advance enables us to examine optimal trade-offs across time and interventions that are more human targeted (i.e, school-based MDA) vs more environmentally oriented (i.e., environmental treatment, EnvTre) when used in isolation and in combination with each other. Investigating these trade-offs using simulation analysis would be a monumental task as the combinatorial nature of the possibilities are significant. Second, we show how the implementation of an optimal environmental treatment reduces the dependency on mass drug administration and that this reduction increases when the planner considers a longer planning period. This latter result highlights potential biases in treatment protocols that are based on simulation analysis using short planning horizons.

1.2 Material and Methods

Our economic–epidemiological model of schistosomiasis captures the realistic situation where a central planning agency needs to decide when, what type, and how much treatment to provide to a remote village where the disease is currently endemic. The objective of the central planning agency is to treat the disease at the lowest possible total cost, where the costs include treatment, transportation, and damages (reduction in human

health). The epidemiological model describes the dynamics of infected intermediate hosts living in the environmental reservoir, the population dynamics of the intermediate host themselves, and the dynamics of the infectious human populations.

The structure of our model incorporates both economic and disease ecological factors that vary based on the nature of the treatments. The model structure and parameters represent schistosomiasis but it is not meant to be a tactical tool. Rather our results are indicative and qualitative. More tactical tools could adopt our optimal control framework but would require adaptations to the particular setting and better data for parameterization.

1.2.1 Model of Disease Transmission

The disease model predicts the dynamics of infection of adult and children in a closed population, and the number of intermediate hosts in the environmental reservoir. Adult and children contract the parasite through contact with the environmental reservoir, which here is a body of water next to the village where the disease is endemic, and contribute parasites in the environment via shedding. We assume that part of the population receives MDA (i.e., the children), meaning that it is impossible to completely interrupt the transmission of pathogens into the environment. The number of infected children and adults can go down over time from natural recovery.

In epidemiology, the basic reproduction ratio R_0 is defined as being the expected number of secondary infections, at a disease-free equilibrium, caused by a typical infected individual over its entire infectious period [25]. In our model, R_0 is a function of relative shedding rates of adult and children, natural recovery rates, and contact rates with the environmental reservoir. New infections of the intermediate hosts depend on the relative shedding rates of adults and children, while the loss of infectious hosts is due to natural mortality and the application of a non-selective environmental treatment that kills both susceptible and infected intermediate hosts. Following Lo et al. [12], intermediate infected hosts cannot reproduce. We model a chemical treatment to reduce the freshwater snails, which are the intermediate hosts of schistosomiasis. Disease model parameters are derived from the

literature on schistosomiasis [6; 7; 9; 11; 12; 16; 17; 20; 26; 27; 28; 29; 30; 31; 32; 33; 34]. See the electronic supplementary material (ESM) for details of the disease transmission model and the parameter levels.

1.2.2 Model of Economic Costs

Economic components of the model include treatment costs, damages, and transportation and management costs. We assume damages are additively separable across children and adults. Treatment of children via MDA often occurs in a school setting [6; 9], which reduces the treatment costs associated with administering the drug to children in the village. Based on this approach, we model the treatment cost of children as the level of MDA treatment times the cost of a dose of praziquantel. The cost of the environmental treatment is linear in the amount of chemical treatment, which assumes realistically that increasing the application either through more chemical per unit area or larger area of application increases the cost in a linear manner. To calibrate the cost, we specify a certain size of environmental reservoir and use estimates for variable costs of snail control (e.g., chemical, labor) from the literature [7; 12; 33].

Damages derive from disability and reduced intellectual function [35] causing lower school participation for children [36] and lower worker productivity for adults [37; 38; 39; 40]. For simplicity, we assume that the per unit damage costs are identical across adults and children for a given infection prevalence, but the cost parameters—representing damages on the whole subpopulation and not just one individual—differ due to the proportions of adults and children in the village. We utilize data from Senegal, a country with a GDP per capita close to Africa's median, to calibrate the proportions of children and adults in our population; this gives us a population composed of 40% children (0-14 years) and 60% adults (15 years and over) [31]. The level of damages are set such that in the absence of treatment, there is a prevalence of 38% in a community of 5,000 people and this yields losses of 550 disability-adjusted life years (DALYs) [12]. The value of a DALY was set to be approximately the median value of the GDP per capita of an African country (approximately $3,000 USD).

We model transportation and management costs as a fixed cost in each period during the planning period regardless of whether treatment is being undertaken. We account for potential economies of scale across the different treatment options with a single fixed cost incurred regardless of whether children or environmental treatments are applied. We parameterize the fixed cost from the literature (see, for example, [6; 7; 9; 11; 12; 32; 33]). See the ESM for details of the cost functions and the calibration of the parameter levels.

1.2.3 Planners Decision

Compared to scenario analysis, optimal control techniques require an assumption about the objective of the planning agency. In scenario analysis, one usually computes, for instance, the average cost of an averted DALY to determine the best policy among the simulated ones; highly cost-effective treatments occur when this average cost is below some threshold (e.g., the per-capita gross domestic product) [11]. With optimal control, we solve for the best (i.e., optimal) policy, conditional on the objective of the planner. In our case, we assume that the objective of the planning agency is to minimize the damages and treatment costs of the disease in a remote village where the disease is currently endemic. The objective function is the net present value of the treatment, damages, and transportation and management costs over a period of years, where we assume a four percent discount rate in the base case.

The main analysis considers a ten-year horizon following the prior literature investigating the cost-effectiveness of schistosomiasis treatment options [4; 12; 41; 42]. We consider longer time horizons in the sensitivity analysis. We also assume that the planning agency does not set any target level in year ten for the level of infection prevalence in humans, the level of infected hosts, and the host population size (specifically, we are allowing free endpoint conditions, which implies a set of transversality conditions in the optimal control problem). This allows us to investigate whether eradication is the cost-minimizing outcome at the end of the horizon rather than imposing it as the solution of the planning agency.

Given the cost functions of MDA treatment for children and of the environmental treatment, the controls appear linearly in the formulation. Solutions to linear optimal control models often have a bang-bang nature. That is, the optimal level of the control resides at one limit (e.g., the maximum) for a period of time then switches off to a singular (i.e., intermediate) level or another limit (e.g., the minimum) in another phase of the solution [43]. In these problems, the optimal solution of the control over time consists of discrete switch times. For example, we might expect that the optimal treatment of children to be at the maximum possible level for a certain period of time and then drop to zero, after the infection level in children drops below some endogenous threshold. Given the literature on the non-compliance with MDA treatments [13; 14; 15], the maximum treatment at any instant is equal to 90 percent of the population of school-age children. The limited effectiveness and compliance of MDA treatment [16; 17] further reduces the extent of successful treatment and transmission reduction.

1.2.4 Analysis

To examine the optimal set of MDA and environmental treatment, we numerically solve the optimal control problem across four different scenarios: no controls, school-based MDA, environmental treatment (i.e., snail control), and school-based MDA and environmental treatment.

We use pseudospectral collocation to solve for the optimal dynamics of treatment and infection over time (see [44; 45; 46] for applications of this technique and see the ESM for more details). We present results from a numerical simulation where initially all state variables are at their no-treatment steady-state levels (sensitivity analyses of initial conditions are presented in the ESM; see Figure A.5). The chosen parameter values imply that without any treatment, the infection prevalence for both the children and adult populations will converge to approximately 38% (consistent with the findings of Lo et al. [12]). The steady-state snail population size will converge to the carrying capacity, while the number of infected snails will converge to 54% of total population.

We investigate the impacts of environmental treatments on school-based MDAs by mapping out how the cost of environmental treatments impact the switch time or time at which the planner stops treating school-aged children. Switching off of MDA earlier represents a reduction in treatments and generally a lower reliance on drug treatment as the primary means to address the disease. We also compare the net present value and its components across the different treatment scenarios, and we normalize to one the value of the no treatment case to make comparison easier between cases.

1.3 Results

At our preferred specification of the parameters, we find that the terminal infection levels of the children population are less than 1% when continuously treated with MDA (Figure 2.A). This continuous treatment differs from a pulse treatment, which occurs on a pre-determined fixed interval. Instead of treating the village, e.g., every year, our control mimics a case where the population is being continuously given MDA. The optimal drug treatment (Figure 2.C) is consistent with previous literature on optimal control of epidemics: the disease needs to be hit as hard as possible and as soon as possible [47]. Environmental treatments alone barely reduce the infection level in the children population (Figure 2.A), while driving the infection prevalence in intermediate hosts to about 2% of total intermediate host population (4% of steady-state infection level; Figure 2.B). Unlike the human infection levels that are driven almost to eradication, there are no damages associated with the infected intermediate host, and the incentive to eliminate the disease in the infected intermediate hosts comes exclusively from its effect in the disease's life cycle and its indirect impact on human populations.

Under the optimal scenario, combining environmental treatments with MDA affects the optimal level of drug administrated to children by reducing the switch time (Figure 2.C). Since an environmental treatment (i.e., EnvTre) reduces the level of contaminated intermediate hosts in the environmental reservoir, the transmission of the disease from the intermediate hosts to human populations is reduced, everything else being equal. As a result, less MDA is needed to fight the NTD. The optimal solution suggests that the level

of environmental treatment is only slightly impacted (reduced time spent at maximum control) with the addition of MDA treatments (Figure 2.D)

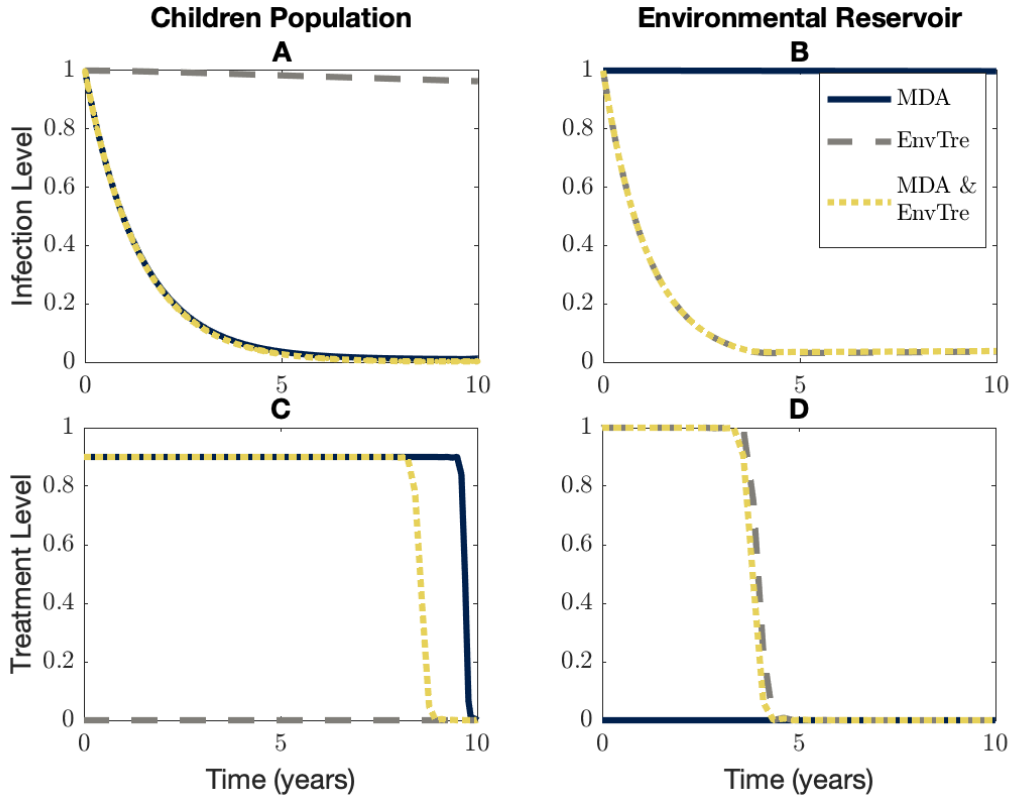

Figure 2: Optimal solutions at base case. This figure shows the change over time of the infection level of the child population (Panel A), the infection level of the intermediate host population (Panel B), the optimal path of drug treatment (Panel C) and the optimal path of environmental treatment (Panel D). Infection levels are expressed as a proportion of their respective steady-state value.

Table 1 summarizes the results in terms of the net present value (NPV, which includes damages and treatment costs), damages (for both the child and adult populations), child MDA costs, costs related to the environmental treatment, and total expenditures (including transportation and management costs) across the different optimal scenarios. By definition, when adding an additional control variable in an optimal control problem, the planner cannot do worse because it could always choose not to utilize this new control variable. To make comparison easier between scenarios, we normalize the measures against the appropriate base (damages are normalized against the no-treatment case).

At these parameter levels, school-based drug treatment (SBDT) reduces NPV by 33 percent. Consistent with Figure 2, environmental treatments only barely impact the level of

infected children and therefore reduce damages by one percent after ten years of optimal chemical snail control. Across both cases, environmental treatments do not contribute to a significant reduction in damages. On the other hand, we find that implementing an environmental treatment reduces the amount of time spent at maximum treatment of MDA by more than one year out of the ten year time horizon, resulting in about a ten percent reduction in MDA expenditures. This cost reduction in MDA could be offset by the increase in costs due to environmental treatments. We find that implementing an optimal environmental treatment requires a slight increase in expenditures, implying that funds are redistributed from MDA to environmental treatments (Table 1; total expenditures are slightly increased when SBDT is combined with environmental treatments). Even though total expenditures are increased, this situation is still preferable given the lower net present value.

EnvTre	**MDA**	**NPV**	**Damages**		**Expenditures**		
			Child	**Adult**	**Child**	**EnvTre**	**Total**
No	**None**	1	1	1	–	–	–
	SBDT	0.67	0.17	1.00	1	–	1
Yes	**None**	0.99	0.98	0.98	–	1	0.99
	SBDT	0.66	0.16	0.98	0.90	0.97	1.08

Table 1: Normalized values of net present value (NPV), damages (reduction in human health), treatment costs (child, environmental treatment (EnvTre), and total which includes transportation and management costs) for when the planning horizon considered by the social planner is $T = 10$ years.

Our results highlight important trade-offs between direct (e.g., treatment of school-aged children) and indirect (e.g., treatment of intermediate hosts) treatments and suggest that the optimal amount of MDA is reduced when the policy is combined with an environmental treatment policy. However, the magnitude of the reduction in MDA due to the implementation of an environmental treatment inherently depends on (i) the costs associated with the environmental treatment (here, the marginal cost of snail control), and (ii) the basic reproduction ratio of the disease, R_0. We find not surprisingly that as the cost of the environmental treatment goes up, the planner reduces the time during which the maximum control is applied (Figure 3, bottom panel). Consistent with our base

case, we find that as the cost of environmental treatment goes down, the planner utilizes less MDA, as measured by the shorter proportion of the time spent at the maximum treatment level (Figure 3, top panel). As the cost increases, we converge to the solution where no environmental treatments is the optimal solution. While we are agnostic on the source of this cost increase, one potential source could stem from the damages of these environmental treatments on other species in the ecosystem (for environmental damages associated with snail control to fight schistosomiasis, see [35; 48; 49; 50]).

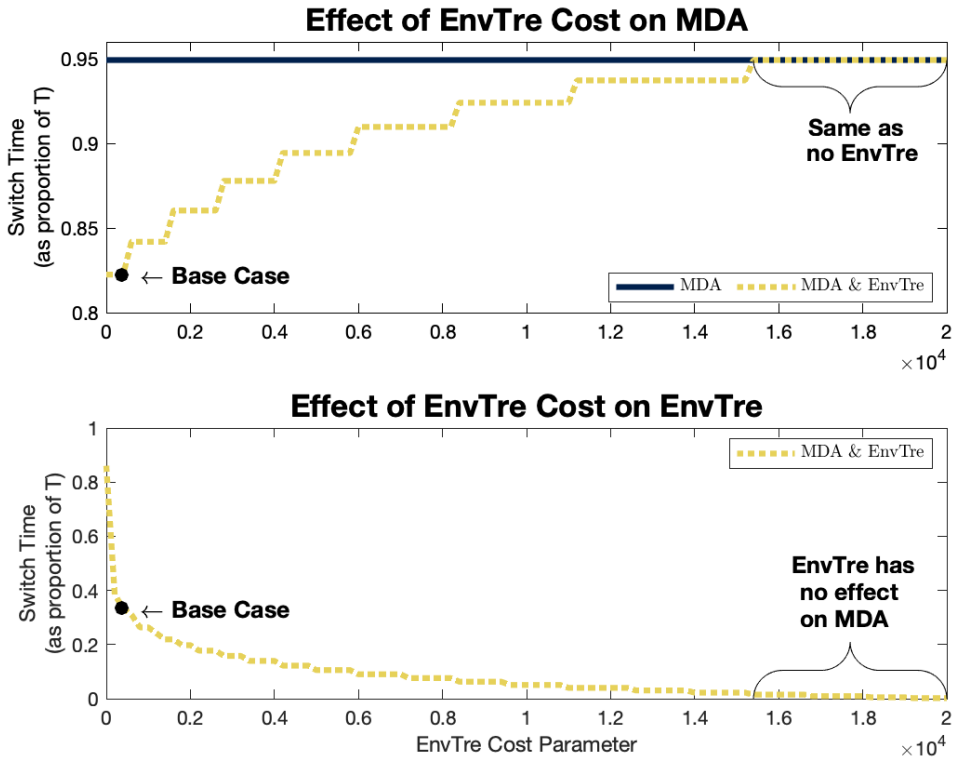

Figure 3: Proportion of total treatment time spent at the maximum level of control (MDA or environmental treatment) as a function of the cost parameter associated with the environmental treatment (EnvTre) . The point given by "Base Case" represents the switch time and environmental treatment cost parameter of Figure 2 and Table 1.

According to Sokolow et al. [20], the expected range of R_0 for schistosomiasis ranges from 1 to 7. While our base case is 3.5, we investigate the range given in Sokolow et al. [20]. There are multiple parameters that affect R_0 (see ESM for derivation) and the ones for which we have less information are the contact rates and the shedding rates. By varying these parameters to vary the R_0, we find that for the majority of the range of R_0, our

finding on the optimal substitution away from MDA to environmental treatment holds qualitatively (Figure A.2, top two panels); the impact seems relatively constant at least between 2 and 7. With a R_0 higher than in our base case, the amount time spent at the maximum level of environmental treatment will be slightly higher, everything else equal, and more so if this higher R_0 is due to higher contact rates, and vice versa (Figure A.2, bottom panel). Only when the R_0 approaches one do we find significant changes to the switch times for MDA and environmental treatment.

Our parameterization of the contact rates (1 infection per 200 water contacts) and shedding rates (1 intermediate host infection per 555 sheds) are consistent with the literature [20; 26], but there are multiple different combinations of these parameters that could yield the same basic reproduction ratio. We investigate potential impacts of these combinations by keeping our base case value of 3.5 constant and varying the level of contact rate relative to the shedding rate. We find that (i) the substitution away from MDA due to the environmental treatment remains approximately the same regardless of the relative levels of the contact rate and shedding rate, and (ii) both MDA and the environmental treatment increase as the contact rate becomes relatively higher in magnitude relative to the shedding rate (Figure A.3)

Following the previous literature [4; 12; 41; 42], we use a ten-year planning period. In our optimal control framework, the implications of a ten-year horizon either imply that costs are no longer incurred after year ten, or that the central planning agency does not consider costs incurred after year ten; both interpretations seem unrealistic. In our model, this implicit assumption explains why, even in a ten year planning horizon, the optimal solution requires an abandonment of MDA. In the prior literature using scenario analysis, the implicit assumption is that treatment will continue indefinitely in the same *ad hoc* pattern. Considering only shorter planning horizons, however, could bias treatment prescriptions to those that work immediately, which might be a good strategy during an outbreak but not necessarily for an area with endemic disease.

To investigate the interaction between optimal treatment prescriptions and planning horizons, we solve the optimal control model over longer time horizons. We do not impose that either treatment must occur after year ten. That is, we could find that the optimal solution is to abandon the village at some point in the future (i.e., both treatments are optimally set to zero). Our results suggest that as the planning horizon increases, the optimal solution is to substitute away from MDA to environmental treatment. For example, while in our base case the environmental treatment reduced MDA switch time by a little over one year, when the planning horizon is 50 years, this reduction is approximately 40 years (Figure 4).

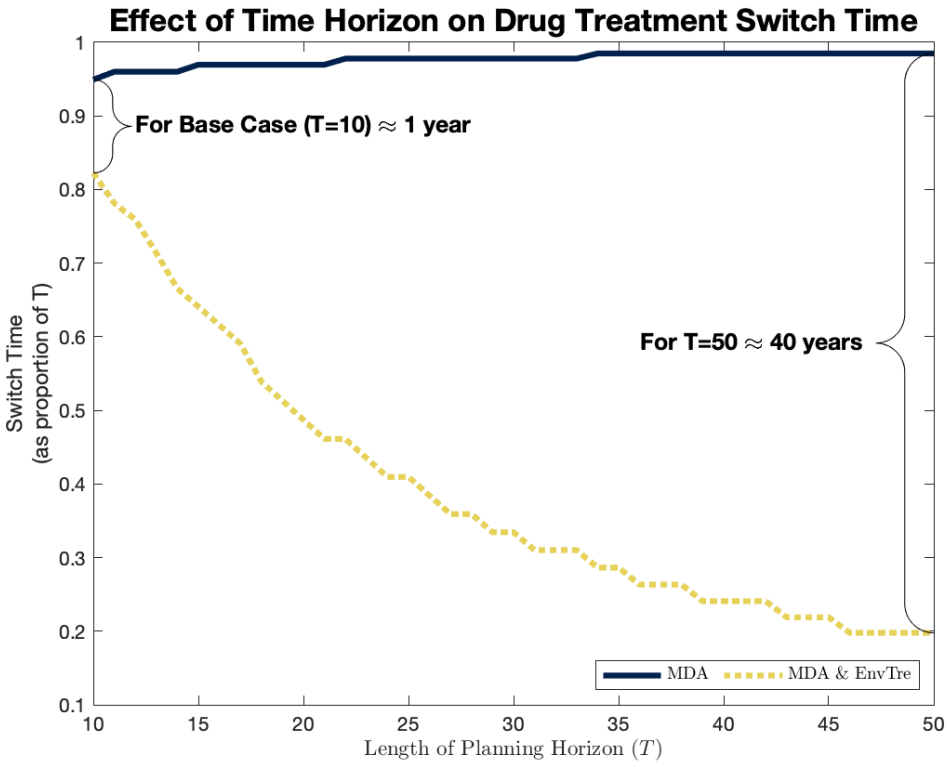

Figure 4: Proportion of total treatment time spent at the maximum level of MDA as a function of the time horizon, T, considered in our analysis. We show the result for the MDA alone case, and for the MDA & EnvTre case.

This reduction in MDA treatment translates into a more than 50 percent reduction in MDA expenditures over the entire planning horizon. As such, implementing optimal environmental treatment does not require significantly more expenditures (see Table 2).

EnvTre	MDA	NPV	Damages		Expenditures		
			Child	Adult	Child	EnvTre	Total
No	None	1	1	1	–	–	–
	SBDT	0.64	0.07	1.00	1	–	1
Yes	None	0.94	0.94	0.94	–	1	0.97
	SBDT	0.59	0.06	0.94	0.47	0.90	1.01

Table 2: Normalized values of net present value (NPV), damages (reduction in human health), treatment costs (child, EnvTre), and total which includes transportation and management costs) for when the planning horizon considered by the social planner is $T = 50$ years.

With a ten year planning period, the optimal treatment went from the maximum to zero, and remained there for the rest of the planning period (see bottom panels of Figure 2). While the same holds for slightly longer planning periods (see for instance Figure A.11 for when $T = 15$ years), we find that this is not necessarily always true. When the planning period is relatively longer, the switch time more often represents the time where treatment goes from the maximum possible level to a non-zero level that varies over time (see Figure A.12 for when $T = 30$ years and Figure A.13 for when $T = 50$ years).

The qualitative nature of our results are robust to several modeling assumptions. If children represent a greater proportion of the total population, everything else equal, the total amount of time spent treating children does not change (Figure A.4). As long as initial levels of infection are at least 20 percent of the no treatment steady-state values, the qualitative nature of the result remains the same; only when initial infection levels approach 10 percent of the no treatment steady-state values do we find a significant reduction in the substitution away from MDA due to the environmental treatment (Figure A.5). For the levels of discounting we considered (0-20 percent), the MDA switch times remain the same (Figure A.6). Because damages are much larger than treatment costs, the discount rate needs to be very high before it has an impact on the MDA switch times. The amount of time spent on the environmental treatment decreases with higher discount rates, because the long-term benefits to environmental treatment are less important to the optimal solution when the discount rate is high (Figure A.6). This latter result is consistent with the findings under longer planning horizons. The MDA switch times are

invariant to the variations in the effectiveness of the environmental treatment (0.6 to 1, base case 0.88; Figure A.9) and to variations of ±50% in the snails' population growth rate (Figure A.10). Only when the population growth rate of snails approaches its death rate that the amount of time spent at maximum treatment reduces.

Our sensitivity analyses reveal that the reduction in switch time of MDA due to the environmental treatment is mainly affected by (i) the value of a DALY (Figure A.7) and (ii) the effectiveness of MDA control (Figure A.8). As the value of a DALY increases, everything else equal, damages due to disease burden become relatively more important than treatment and transportation costs; to compensate for the relative increase in damages, optimal MDA treatment needs to last for a longer period of time. Hence, with a higher value of a DALY, the substitution away from MDA to the environmental treatment reduces (Figure A.7) because higher MDA effort is preferable, everything else equal. We assumed in our main analysis that treatment was effective 80 percent of the time. Our sensitivity analyses reveal that by improving the effectiveness of drug treatment, the substitution away from MDA to the environmental treatment could be significantly more important (Figure A.8). Higher MDA effectiveness reduces the amount of time spent on MDA treatment, and even more so when combined with an environmental treatment due to reduced reinfection.

1.4 Conclusion

We show the potential value of using integrated treatment guidelines. We find that combining environmental controls and mass drug administrations (MDAs) can significantly reduce the time span over which one has to administer drug treatment, especially when considering a long-term planning horizon. Although WHO recognizes both the advantages [51] and the cost-effectiveness [52; 53] of environmental treatments (in particular snail control), its priority is on MDAs since the development of an anthelmintic drug, praziquantel. School-based deliveries in particular are now the main focus of WHO [6] given the facility to deliver treatment in schools and that children are usually associated with higher disease burden [54].

However, few studies have demonstrated the optimal distribution of integrated approaches and under what conditions different treatments should be used in combination or in series. Specifically, we show that, to achieve an optimal outcome in terms of minimized costs and damages, MDA usage rates can be reduced when used in combination with environmental controls. Similarly, using guidelines that are independently optimal but jointly non-optimal for MDA and environmental controls might lead to inefficiencies: excessive usage of public funds and over-utilization of drug and environmental treatments. In our analysis, we assume unlimited public funds and perfect flexibility of these funds across time, relevant future work could investigate the role of budget constraints and lack of flexibility of rolling funds over from year to year or from one type of treatment to another.

We also show that when transportation and management costs of different types of treatment can be combined in one coordinated program, utilizing both types of controls, instead of only using one control strategy (MDA or environmental control), does not significantly increase total expenditure over 10 to 50 year time horizons. If these costs were only present during the treatment period, then the reduction in time spent treating due to using the combination of multiple types of treatment could lead to further reduction in costs. These additional savings would reinforce the importance of considering an integrated approach to using both drug and environmental treatments.

The environmental treatment we consider in this paper consists of a chemical treatment of the environment. Such a treatment may have a declining efficacy over time. For instance, chemical pesticides used against mosquitoes in malaria-endemic areas have faced limitations due to resistance evolution, non-target effects, and environmental damage [55]. For schistosomiasis, it is well documented that molluscicide niclosamide (the chemical compound used in snail control) can be toxic to other species [35; 48; 49; 50]. Future work could include both the potential ecosystem damages from environmental treatment and potential reductions in efficacy over time. Another possible path is to investigate the feasibility of interventions that focus on reducing pathogen prevalence in the environment that might not have these additional damages or issues with declining efficacy.

For schistosomiasis, there is recent evidence in support of an ecological intervention where snail predator populations are restored [20]. Biological controls using snail predators (e.g., fish, prawns, ducks, crayfish) aid in schistosomiasis control as they reduce snail-to-man transmission by feeding off of the intermediate host population [20]. A potential ancillary benefit of introducing this treatment is the support of fisheries and aquaculture revenue, since many of the candidate natural enemies of snails are also seafood commodities. In fact, this might be a case where treatment does not only improve health outcomes directly but indirectly offers a source of sustainable development that could address food insecurity [42]. However, aligning the incentives of those who indirectly benefit from aquaculture or fisheries restoration with the public health costs associated with schistosomiasis could be challenging [56]. Optimal control methodologies, like the one applied here, are a fruitful approach to understanding the potential benefits and costs of aquatic snail predator restoration or aquaculture for reductions in disease burden and sustainable development.

Building on our results that redistributing funds across controls (e.g., from MDA to an environmental treatment) can be cost-effective, another important area for future research is also considering the optimal gains from redistributing funds across diseases (e.g., from HIV to NTDs, see Mbah et al., 2013) [57].

2 Chapter 2: "Spatial Allocation of Scarce COVID-19 Vaccines"

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 vaccines based, for example, on population size, in the interest of speed and workability. Utilizing economic–epidemiological modeling, we benchmark the performance of this *ad hoc* allocation rule by comparing it to the rule that minimizes the economic damages and expenditures over time, including a penalty cost representing the social costs of deviating from *ad hoc* allocations that favor speed and workability. Under different levels of vaccine scarcity and different demographic characteristics, we consider scenarios where length of immunity and compliance to travel restrictions vary, and consider the robustness of the rules when assumptions regarding these factors are incorrect. The benefits from deviating are especially high when immunity is permanent, when there is compliance to travel restrictions, when the supply of vaccine is low, and when there is heterogeneity in demographic characteristics. Interestingly, a lack of compliance to travel restrictions pushes the optimal allocations of vaccine towards the *ad hoc* and improves the relative robustness of the *ad hoc* 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.

2.1 Introduction

Now that several vaccines against coronavirus disease 2019 (COVID-19) have been developed, an ongoing question for policymakers around the globe is to determine how to allocate the limited supplies. Most of the scientific literature on allocation has fo-

cused on demographic considerations within one jurisdiction [58; 59; 60] or on a global scale [61; 62; 63; 64]. This prior work has made important contributions to the debate. A missing piece in the allocation question is how to divide up limited quantities across jurisdictions (e.g. state, counties) that might have different demographic and epidemiological characteristics. A report on the allocation of a COVID-19 vaccine by the U.S. National Academies of Sciences, Engineering, and Medicine (NASEM) [65] 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 a fixed *ad hoc* allocation rule by comparing it to the optimal rule conditional on the level of scarcity of the vaccine. Throughout this paper when we refer to the "*ad hoc* allocation," what we mean is a rule of thumb that favors "speed and workability," so we follow the U.S. NASEM [65] allocation recommendation based on the jurisdictions' population size. The optimal rule we consider is assumed to be one that minimizes the economic costs from health-related damages, vaccine expenditures, and a workability cost imposed on the planner for deviating from the *ad hoc* rule.

In a world where two jurisdictions are identical in terms of population, the *ad hoc* rule would divide the limited supply equally between the jurisdictions. However, it is much more likely that two jurisdictions, even if equally sized, have heterogeneous levels of infections (e.g. in terms of cases) at the time a vaccine is licensed and starts to be administered. Based on prior literature on spatial-dynamics of disease management, heterogeneity in infection levels may lead to significant deviations between the optimal spatial allocation and the *ad hoc* rule (see [66] for example).

Mechanisms leading to heterogeneous infection include the timing of the outbreak, demographic characteristics of the population (e.g. age structure [67] and essential worker status [68]), and the implementation of and compliance with preventative interventions;

see [69; 70] for more details on how SARS-CoV-2 (i.e. the virus that causes COVID-19) prevalence varies across space. While compliance to preventive measures may seem independent from vaccine allocation, it affects the initial conditions (i.e. the conditions before the vaccine is licensed and starts to be administered) 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 movement of the virus from one jurisdiction to another. When regions are non-interacting, Brandeau et al. [71] 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. When regions are interacting, Rowthorn et al. [47] show when there is no immunity (i.e. in an SIS model) that treatment should be preferentially directed towards the region that has the lower level of infection. While these results indicate that a fixed *ad hoc* rule is less cost-effective in an SIS model, whether compliance to travel restrictions makes the *ad hoc* rule relatively more cost-effective in the case of COVID-19 is an open question.

Our findings illustrate that the vaccines should be optimally allocated over time depending on: (i) if the jurisdiction has initially a lower or higher disease burden, (ii) if immunity is permanent (see Zhou et al. [72]) or temporary (Gersovitz and Hammer [73] already pointed out that the optimal allocation is conditional on the duration of immunity), (iii) whether there is compliance to travel restrictions or not, (iv) the amount of vaccine available, and (v) the average demographic characteristics of the population (i.e. age structure and essential worker status). We proxy variability in demographics by assuming that the population of one jurisdiction has a higher case-fatality ratio (mimicking an older population) or a higher contact rate (mimicking a population containing more essential workers) than the other. We find that the benefits of deviating from the *ad hoc* rule are especially high when immunity is permanent, when there is compliance to travel restrictions, when the vaccine supply is low, and when there is heterogeneity in demographic characteristics. Allocating a vaccine based on an *ad hoc* allocation rule generally leads to an over-utilization in jurisdictions where disease prevalence is higher, an under-

utilization in jurisdictions where disease prevalence is lower, and overall a higher number of cumulative cases. Whether these inefficiencies outweigh the "speed and workability" inherent in *ad hoc* rules is an important question for policymakers. Our research can aid in that discussion by illuminating the trade-offs involved in such complex epidemiological, economic, and social decisions by providing optimal benchmarks from which to compare *ad hoc* rules.

While the optimal allocation is conditional on a number of factors mentioned above, the science remains unresolved on the duration of immunity to SARS-CoV-2, and it is difficult to anticipate and subsequently estimate the extent to which populations in different jurisdictions comply with the travel restrictions. On the other hand, the *ad hoc* allocations have the advantage of being based on easily observable factors (e.g. a jurisdiction's population size). To gain insights into the robustness of optimal and *ad hoc* policies in the presence of such uncertainties, we investigate the economic and public health consequences that could occur if we design an optimal policy or evaluate the performance of *ad hoc* rules under a set of assumptions on immunity and compliance that turn out to be incorrect.

We make a number of contributions to the literature. First, we develop an economic–epidemiological model and solve for the optimal allocation of vaccines over time to minimize the economic costs from damages, vaccine expenditures, and a workability cost imposed on the planner for deviating from the *ad hoc* rule. Prior literature considering the trade-offs involved with *ad hoc* rules does not consider that deviating from them entails potential workability costs (see, for example, [74]). Second, we consider how vaccine allocations are influenced by compliance with preventative interventions (i.e. travel restrictions). Third, we demonstrate how vaccine allocations are dependent on various demographics (i.e. age structure and essential worker status). Fourth, we show that, in general, optimal rules are robust to incorrect assumptions about the duration of immunity but differences in public health outcomes (cumulative cases) appear when compliance to a travel restrictions is assumed when in fact there is not compliance; it is, however, much preferable from a public health outcome perspective to comply with travel restrictions.

The paper is divided as follow. In Section 2.2, 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 2.3 presents the results while Section 2.4 concludes the paper.

2.2 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 to allocate to two jurisdictions where disease burden is heterogeneous at the moment the vaccine is licensed and starts to be administered. We assume that the objective of the central planner is to minimize costs across both jurisdictions, including damages associated with the morbidity and deaths of infected individuals, the expenditures related to the pharmaceutical intervention, and a penalty cost mimicking the increased workability costs incurred for any deviation from the *ad hoc* allocation. The dynamics of SARS-CoV-2 are modeled using an SEIR epidemiological model, which tracks the change over time of the susceptible (S), exposed (E), infected (I), and recovered (R) populations for two separate jurisdictions (see Appendix B.1 for more details on the calibration of the model). We note that while we generally talk about these jurisdictions as being two different states, they can very well represent two counties, or regions within one state.

2.2.1 Modelling Different Types of Intervention

There are two different types of interventions we consider: travel restrictions and vaccines. We assume that travel restrictions affect both jurisdictions simultaneously (e.g. by an order from the central government), and that the populations either comply perfectly or imperfectly to the travel restrictions (for examples of optimal lockdown policies see, e.g., [75; 76]). When compliance is perfect, individuals in different jurisdictions do not interact with each other and thus susceptible individuals can only get infected by being in contact with some infected individual in their own jurisdiction. When compliance is imperfect,

susceptible individuals from one jurisdiction can also travel to the other jurisdiction where they can be in contact with infected individuals, or infected individuals from one jurisdiction can travel to the other jurisdiction and infect susceptible individuals there; this discrete shift in the number of contacts effectively increases the transmissibility of the virus (see Appendix B.1 more details).

We assume that the analysis starts when a vaccine has already been developed, licensed, and is available in a relatively high quantity. For simplicity, the amount of available vaccine is assumed to be exogenous to the model and fixed over time, which is likely given the short time frames we consider in the paper. However, we consider different levels of vaccine supply to investigate how different levels of vaccine scarcity may affect their optimal allocation. In our model, vaccines reduce the pool of susceptible individuals by providing them with immunity from the virus, as early evidence suggests that vaccines could be transmission blocking in addition to preventing severe disease [77].

2.2.2 Model of Disease Transmission

We use a frequency-dependent [78] 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 that is either susceptible, exposed, infected, or recovered (see Figure 5). We also consider scenarios where immunity is temporary (i.e. lasts 6 months, for more details see [79]), thus also using an SEIR–Susceptible (SEIRS) model (for COVID-19 applications see, e.g., [80; 81; 82; 83; 84; 85; 86; 87; 88; 89; 90]). In such scenarios, the R_i recovered individuals are immune for a mean period of $\frac{1}{\omega}$ months.

In each jurisdiction i, the S_i susceptible individuals are in contact with the I_i infected individuals of their own jurisdiction at a rate of β_{ii} and are in contact with the I_j infected individuals of the other jurisdiction at a rate of β_{ij}. We assume $\beta_{ij} = 0$ (i.e. no mixing between jurisdictions) when there is perfect compliance to travel restrictions, and $\beta_{ij} > 0$ if not. To highlight the role of travel restriction compliance and initial disease burden, we initially assume that the contact rate is identical across jurisdictions, meaning that

$\beta_{11} = \beta_{22} = \beta_{ii}$ and $\beta_{12} = \beta_{21} = \beta_{ij}$ (in Section 2.3.2 we relax this assumption and investigate the optimal allocation when there is heterogeneity in the contact rate). We assume there is no permanent migration of individuals from one jurisdiction to another (see for instance [91] and see [92] for an example applied to COVID-19) in the sense that individuals who do not comply with travel restrictions do not permanently move to the other state, but instead travel to it temporarily. An implication is that we are assuming that the two jurisdictions are close enough for such travel and mixing to be economically feasible.

We model the control variables for vaccines as non-proportional controls, i.e. available in a constant amount each month [47; 60; 93; 94]. 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 re-vaccinated in subsequent months.

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

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

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

in Section 2.3.2. The growth of the infected individuals is

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

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

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

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

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

In keeping with much of the previous economic epidemiology literature [73] as well as recent applications to COVID-19 (see for example [95]), we have omitted natural births and non-COVID-related deaths due to the short time frame of our model (4 months) and assume reductions in international travel [96] effectively lead to a closed population (i.e. there is no exogenous importation of infected individuals). See Appendix B.1 for more details about the parameterization of the epidemiological model.

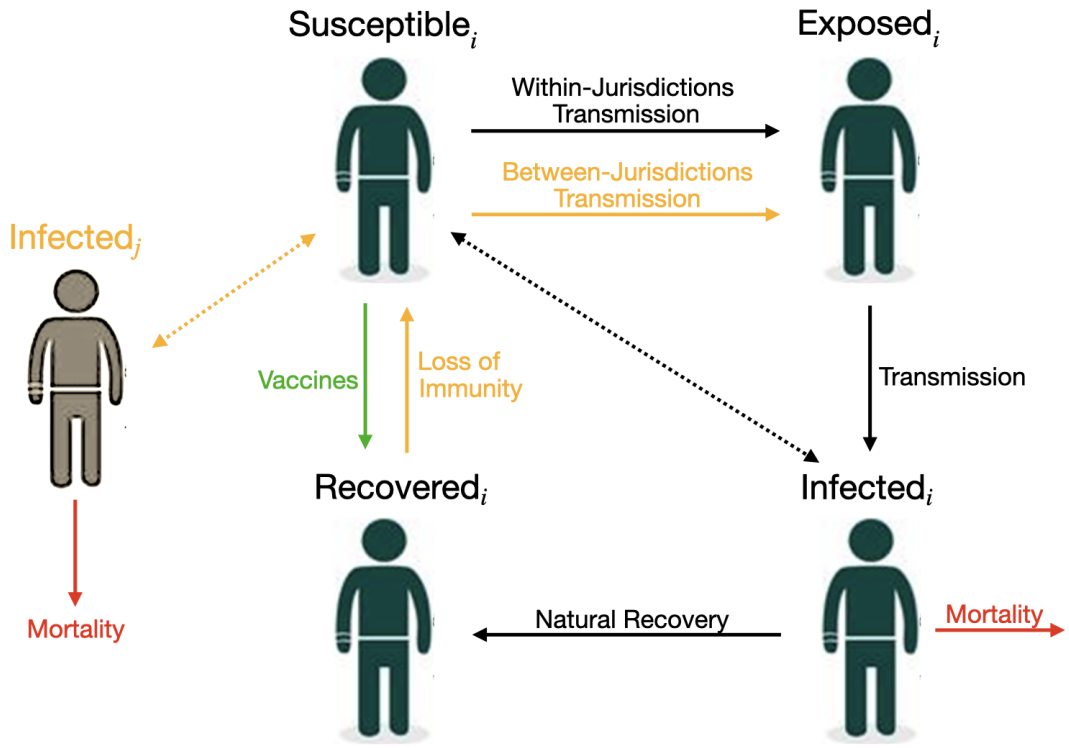

Figure 5: This schematic shows the model interventions and disease transmission pathways for our model of COVID-19. The full lines represent the transition between, or out of, compartments while the dotted lines represent contact between susceptible and infected individuals. Black lines represent situations that do not vary, while yellow lines represent key factors that we vary in our model to see how they impact our results. The green line represent the vaccines and the red line represents mortality.

2.2.3 Modelling *Ad Hoc* Allocations

We model an *ad hoc* allocation rule that favors "speed and workability" [65]. We follow the NASEM approach [65] 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.

In the *ad hoc* scenarios, we model the allocation rule as an inequality because towards the end of the horizon after periods of vaccinations, the level of susceptible in the population

may be such that the limited supply of vaccines is not an issue. Other *ad hoc* rules are possible, such as, allocate all to the largest or smallest population [74], but we concentrate on the one currently being advocated for by NASEM [65].

2.2.4 Model of Economic Costs

The model of economic costs include damages related to morbidity and deaths, costs spent on the vaccines, and the workability cost described above that is incurred for any deviation from the *ad hoc* allocation rule. Damages represent consequences related to a temporary disability associated with severe or critical symptoms, and loss of life in the worst cases. The damages are assumed to be linear and additively separable across jurisdictions, meaning that they are identical across individuals and across jurisdictions. The marginal value of damages (i.e. the damages associated with the death of one individual) is assumed to be constant over time and given by the value of a statistical life (VSL) that the U.S. Environmental Protection Agency [97] uses (see Appendix B.1 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 [98]), we use the disability value associated with lower respiratory tract infections. The damage function for Jurisdiction i is

$$c_i(I_i) = (w + \varphi_i)cI_i \tag{7}$$

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

We model a scenario where the central planner is focused on the allocation of vaccines where the costs for its development have already been incurred. This implies that vaccine development costs have already been utilized (in technical terms we say that the costs are sunk) and therefore do not affect the decision of the central planning agency. We model the vaccination cost as linear, where the cost parameter represents the cost of treating

one individual. The vaccine cost function is denoted $c_{V_i}(u_{V_i})$, with $i = 1, 2$. We assume that the vaccination cost is additively separable across jurisdictions such that we denote the cost of treating u_{V_i} individuals as

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

where c_V represents the cost of treating one individual via vaccine. Calibration of the cost parameter is based on current vaccine prices (see Appendix B.1 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 [99]). The workability cost function is:

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

where c_A is the parameter associated with the workability cost. When the gains from deviating from the *ad hoc* allocations (i.e. a reduction in damages in one jurisdiction) outweigh the costs (i.e. an increase in damages in the other jurisdiction and the increased workability costs incurred), the central planning agency will prioritize this allocation as it will lead to lower total costs. By imposing the *ad hoc* rule *ex ante*, the decision-maker is essentially assuming that this workability cost is infinite. Everything else being equal, we expect that the presence of the workability cost will push the optimal allocation towards the *ad hoc* rules (see Figure B.18 for a sensitivity analysis of our results to the workability cost parameter). Therefore, when we do find deviations, we need to consider that these include this workability cost and if workability costs smaller, then the deviations and trade-offs would be greater.

2.2.5 Planner's Objective

In optimal control theory, the best, or optimal, path of the control variables (here the allocation of the limited supply of vaccines) is conditional on the objective of the central planning agency. We assume that the objective is to minimize the economic damages and the costs of the pharmaceutical intervention across jurisdictions over time, rather than a solely epidemiological objective (see for instance [47]). The objective function is the net present value of damages, expenditures related to vaccination, and the workability cost over an exogenously determined planning horizon (4 months). Specifically, the planner's objective is:

$$\min_{u_{V_1}, u_{V_2}} \int_0^T e^{-rt} \Big\{ c_1(I_1) + c_2(I_2) + c_{V_1}(u_{V_1}) + c_{V_2}(u_{V_2}) + c_A(u_{V_1}, u_{V_2}, N_1, N_2) \Big\} \mathrm{d}t \qquad (10)$$

where r is monthly discount rate. The planner solves equation (10) over a fixed time interval, T, subject to equations (1), (2), (3), (4), (5), along with constraints on availability of vaccines ($u_{V_1} + u_{V_2} \leq \bar{u}_V$), non-negativity conditions, physical constraints on vaccines, initial disease burdens in each jurisdictions, and free endpoints (see discussion on terminal conditions in the next section). In the *ad hoc* scenarios, we also impose equation (6).

2.2.6 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 is already licensed and starts to be administered) is calibrated using the epidemiological model (equations (1), (2), (3), (4), and (5)). At the beginning of the outbreak, we assume that, in each jurisdiction, there is one exposed individual in an otherwise entirely susceptible population of 10 million individuals (approximately the population of Michigan), 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.2 for more details. In Section 2.3.2 when we consider heterogeneity in demographic characteristics (varying case-fatality

ratio and contact rate), we modify the initial conditions accordingly assuming an identical timing in the outbreak of the disease.

We impose no conditions on the number of susceptible, exposed, infected, and recovered individuals at the end of the planning horizon; in technical terms, we say that the state variables are free (see Appendix B.2 for more details). Under our free endpoint conditions, there is a transversality condition (i.e. a necessary condition for the vaccine allocation to be optimal) for each state variable that requires the product of the state variable (S_i, E_i, I_i, R_i or N_i) and its corresponding costate variable (i.e. the shadow value, or cost, associated with the state variable) is equal to zero. Hence, at the end of the time horizon, either the state variable equals zero, the shadow value associated with the state variable equals zero, or both. In any case, allowing state variables to be free guarantees that the terminal levels of the state variables are optimally determined. Another possible assumption could be that over a fixed interval we find the optimal policy such that at the end of the horizon there is a given percent reduction in infected or susceptible individuals. Our approach nests this more restricted scenario.

2.3 Results

To examine the optimal allocations of vaccine over time, we numerically solve the optimal control problem across three different scenarios: no controls, optimal vaccine allocation, and *ad hoc* vaccine allocation. We investigate how to allocate vaccines by mapping out the different allocation rules for different immunity–travel restrictions–capacity scenarios. Any deviation from the *ad hoc* allocation rule is optimal despite incurring the workability cost. As the workability cost parameter c_A goes to zero, the problem becomes linear in the controls where the optimal allocations in linear problems follow singular solutions. We use pseudospectral collocation to solve for the optimal dynamics of vaccine and infection over time, which converts the continuous time optimal control problem into a constrained nonlinear programming problem solving for the coefficients of the approximating polynomials at the collocation nodes (see [45; 100] for other applications, and see Appendix B.2 for more details on this technique).

We present the results for our preferred specification of the parameters (i.e. following what was estimated in the literature; see details in Appendix B.1) and for the case where immunity is permanent and the case where immunity is temporary. We detail the optimal deviation based on whether the populations of the different jurisdictions are compliant to travel restrictions or not, and for different levels of capacity constraints. The total available quantity of vaccine in a given time period (i.e. a month; $\bar{u}_V$) is based on a certain percentage (5%, 10%, or 15%) of the total population size. We focus our analysis on the period of time when the scarcity of the vaccine constraint is binding, as once the constraint relaxes the allocation question becomes moot.

2.3.1 Base Case: Homogeneous Demographic Characteristics

Compliance to travel restrictions impacts the optimal allocation of vaccines, regardless of whether immunity is temporary or permanent and regardless of the amount of vaccine available. Noncompliance to travel restrictions reduces both the oscillation (i.e. back-and-forth movement of resources between jurisdictions) of the optimal allocation and the amplitude of the deviations from the *ad hoc* rule (see Figure 6 for when immunity is permanent, and see Figure B.1 for when immunity is temporary). Because noncompliance to travel restrictions decreases the structural heterogeneity in the system, the optimal allocation of vaccine converges towards the *ad hoc* allocation when populations mix with each other. This result clearly demonstrates how the performance of the allocation rule is dependent on how citizens in the jurisdictions comply with travel restrictions.

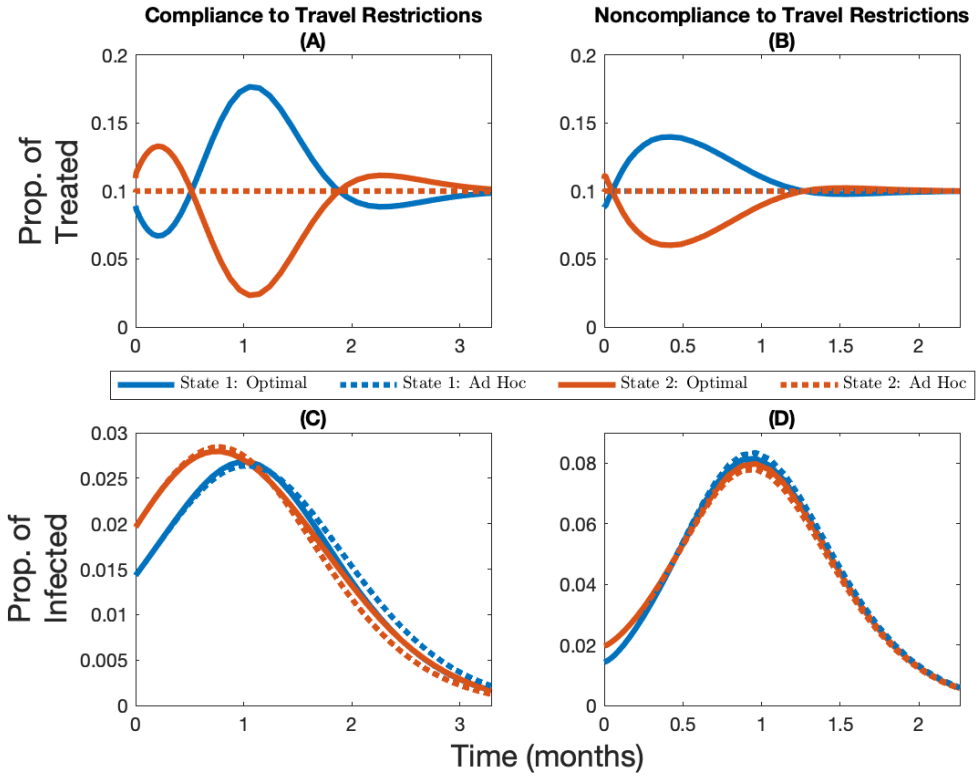

Figure 6: **Vaccine allocation with and without compliance to travel restrictions**. Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether there is compliance to travel restrictions (panels A and C) or not (panels B and D) for the case where the vaccine capacity constraint is 10% and immunity is permanent. Note the changing y-axis in panels C and D in order to better highlight the infection levels.

Noncompliance to travel restrictions leads to the initially less infected state being favored by the optimal allocation for low levels of vaccine capacity (e.g. 5% capacity; see Figure 7 Panel A for when immunity is permanent and Figure B.4 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 7 panels B and C for when immunity is permanent and Figure B.4 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 should be given to the more infected state at

the beginning of the time horizon (see figures 7 and B.2 for the case where immunity is permanent; see figures B.3 and B.4 for the case where immunity is temporary).

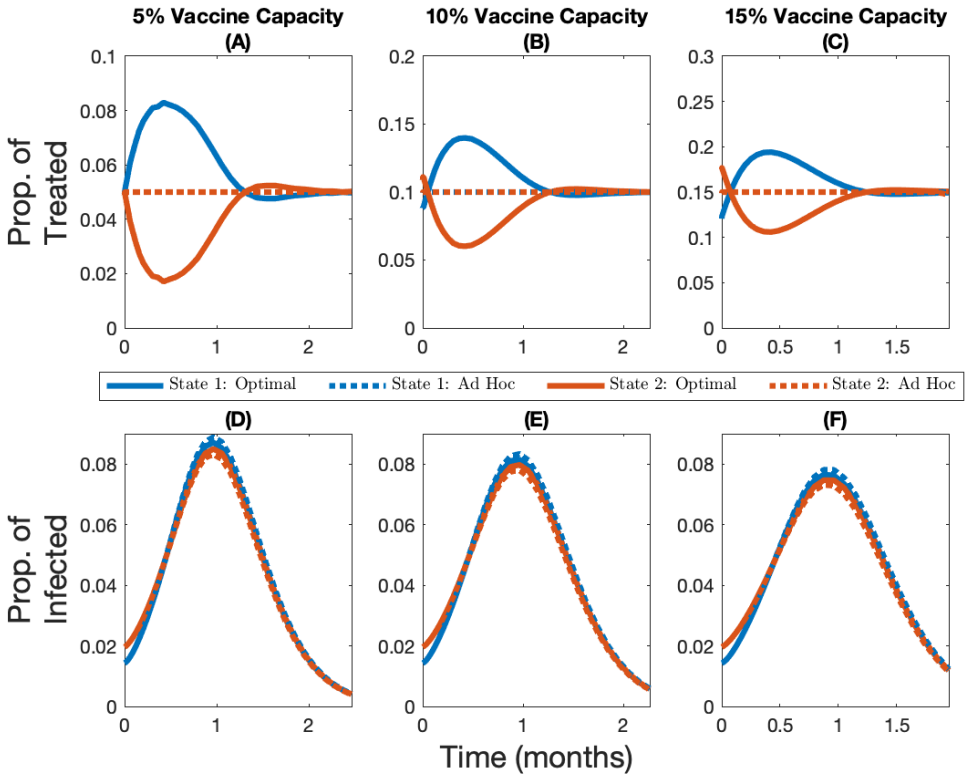

Figure 7: **Vaccine allocations under different levels of scarcity without compliance to travel restrictions**. Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether capacity is 5% (panels A and D), 10% (panels B and E), or 15% (panels C and F), for the case where immunity is permanent and there is no compliance to travel restrictions.

Interestingly, temporary immunity has a different effect on the optimal vaccine allocation depending on whether or not populations are compliant to travel restrictions. When populations comply with travel restrictions, temporary immunity increases the oscillation of the optimal allocation because benefits from vaccination are only temporary, and since the population gradually loses its immunity, it forces more back-and-forth movement of resources between jurisdictions (see Figure B.5). When populations do not comply with travel restrictions, temporary immunity reduces the amplitude of the deviations from the *ad hoc* rule because it further dampens the structural heterogeneity in the system, since the infection and recovery level of both jurisdictions will eventually reach the same

positive steady-state level (recall the only heterogeneity in the system is the initial disease burden in the base case).

While the optimal allocation of vaccine is unequal from a resource allocation perspective, it equalizes the current infection levels across jurisdictions (Figure 6 Panel C). As the vaccine capacity increases, however, the *ad hoc* allocation rule performs better and in turn the amplitude of the optimal deviation decreases (see Figure B.2 panels A, B and C, or Figure B.3 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 8). Hence, the optimal allocation makes the current infection level more equal, while the *ad hoc* allocation makes cumulative infection more equal. In fact, in all scenarios considered, the optimal allocation will lead to lower cumulative damages in the less infected jurisdiction but higher cumulative damages in the most infected jurisdiction (see Figure 8 with vaccine capacity of 10%, and see figures B.7 and B.8 with vaccine capacity of 5% and 15%, respectively).

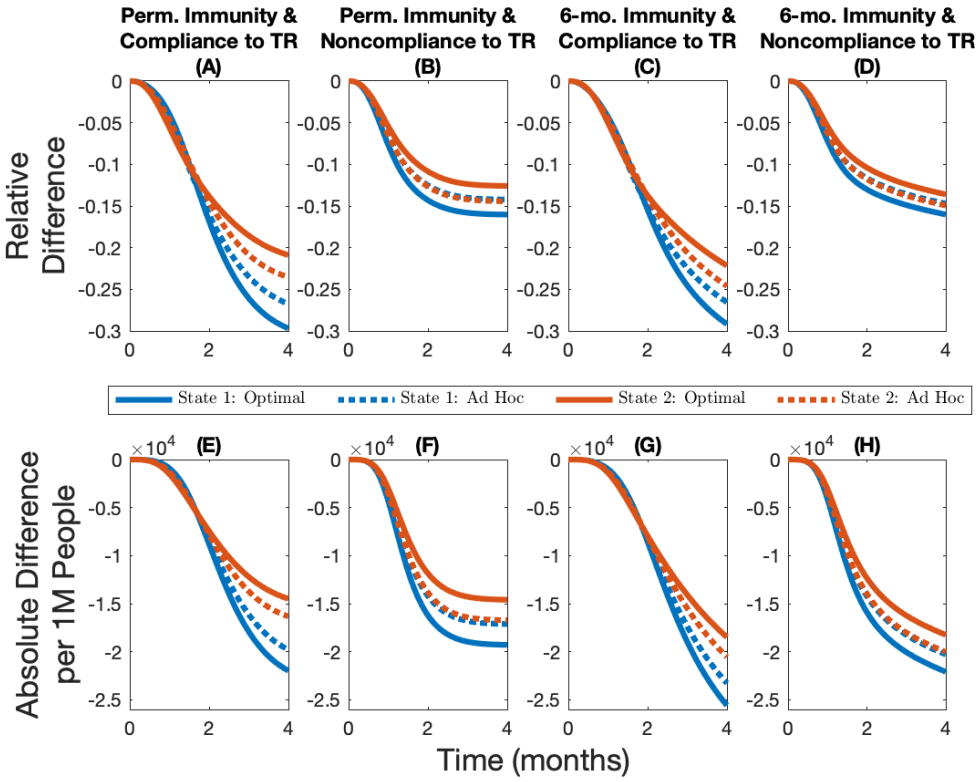

Figure 8: **Epidemiological outcomes under different scenarios with vaccines.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-vaccine case for different immunity–travel restrictions scenarios and for when vaccine capacity is 10%.

2.3.2 Heterogeneous Demographic Characteristics

In our main analyses, we introduced heterogeneity in infection across jurisdictions by assuming a different timing in the outbreak of the virus. While this leads to heterogeneity in the disease burden, there are other mechanisms that could lead to similar differences in the disease burden at the time the vaccine is licensed and starts to be administered. For instance, in jurisdictions that have an older population on average, we expect SARS-Cov-2 to have a higher case-fatality ratio [67] which would effectively decrease the transmissibility of the virus (see Appendix B.1 more details). If we assume that one jurisdiction has a higher case-fatality ratio and start the initial outbreak at the same time, then we find the population with the highest case-fatality ratio also has the largest population of susceptible individuals at the time the vaccine is administered. The heterogeneity in initial disease burden stems from the lower transmissibility in the

jurisdiction with the higher case fatality ratio. The heterogeneity in case fatality ratio therefore not only leads to heterogeneity in the initial disease burden but it also implies that the benefits of vaccination are no longer homogeneous across jurisdictions. These differences lead the optimal allocation to favor even more the least burdened jurisdiction, which is also the most vulnerable (aka more older individuals) of the two populations. Overall, introducing heterogeneity in case fatality ratio strengthens our main set of results (see Figure B.9 for when immunity is permanent and see Figure B.10 for when immunity lasts 6 months).

Another source of heterogeneity in infection could stem from one jurisdiction having more essential workers than the other (for more details on how the risk of infection is occupation-dependent, see [68]). In our model, we can capture this by considering spatial heterogeneity in the contact rate, where a higher contact rate proxies for more essential workers. This in turn leads to a higher initial disease burden in the jurisdiction with higher contact rate. We find that priority is given to the state with a higher contact rate (aka more essential workers) in almost all cases. As the state with a higher contact rate gets vaccinated, we eventually shift priority to the state with the lowest contact rate because either the number of cases starts decreasing in the jurisdiction that has the higher contact rate, or the low-contact state eventually reaches a point where its infection level becomes higher than the high-contact state (see Figure B.11 for when immunity permanent and see Figure B.12 for when immunity lasts 6 months). In the case when immunity lasts 6 months and there is compliance to travel restrictions (Figure B.12, Panel A), we find that priority is given to the low-contact jurisdiction, which contradicts many notions of fairness associated with vaccine allocation. With the gains of vaccination temporary and no movement of people, it turns out that the greatest return per vaccination is in the place where you can best avoid future cases (low contact rate jurisdiction). This prioritization is only fleeting however and there is more back-and-forth movement of resources between jurisdictions in this case, even though the workability cost is being incurred each time.

2.3.3 Robustness of Spatial Allocations

There is significant uncertainty associated with the duration of immunity (i.e. if it is permanent or temporary) and to what extent populations comply with travel restrictions. One argument for the *ad hoc* allocation is that uncertainty in these parameters makes the optimal allocation impossible to achieve. This uncertainty is not yet resolved and public health officials have to choose vaccine allocations based on potentially incorrect assumptions. We compare the robustness of the optimal spatial allocation to the *ad hoc* allocation. By definition, the optimal allocation minimizes the net present value of the health-related damages and total expenditures (including vaccine expenditures and the workability cost incurred because of the deviations from the *ad hoc* allocation), and thus cannot do worse on this dimension than the *ad hoc* allocation. We measure robustness by first inserting the optimal solution under one set of assumptions into the disease dynamics under another set and compute the changes in total expenditures (i.e. the pharmaceutical intervention and the workability cost) and public health outcomes (cumulative cases) over time. We then calculate the distance of these changes in percentage terms to the optimal solution derived under the "correct" assumptions (represented by the point $(0, 0)$ in Figure 9). For example, suppose immunity is permanent and there is perfect compliance to a travel restriction. We derive the optimal policy under these assumptions and use it to measure the robustness of the optimal policies that are derived under assumptions that immunity is temporary and/or there is noncompliance. The *ad hoc* policies being based on observable factors are then compared to the incorrectly applied optimal policies. We illustrate the case for 10% scarcity and include other scarcity cases in Appendix B.3.

When demographic characteristics are homogeneous across jurisdictions, we find overall that immunity length has a lesser impact on both economic and epidemiological outcomes than compliance to travel restrictions (compare the distance from the origin between the plusses and the stars in Figure 9). There are more nuanced trade-offs, however (e.g. compare position of the stars across the panels in Figure 9). 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 9 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 B.13 and B.14 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 B.14).

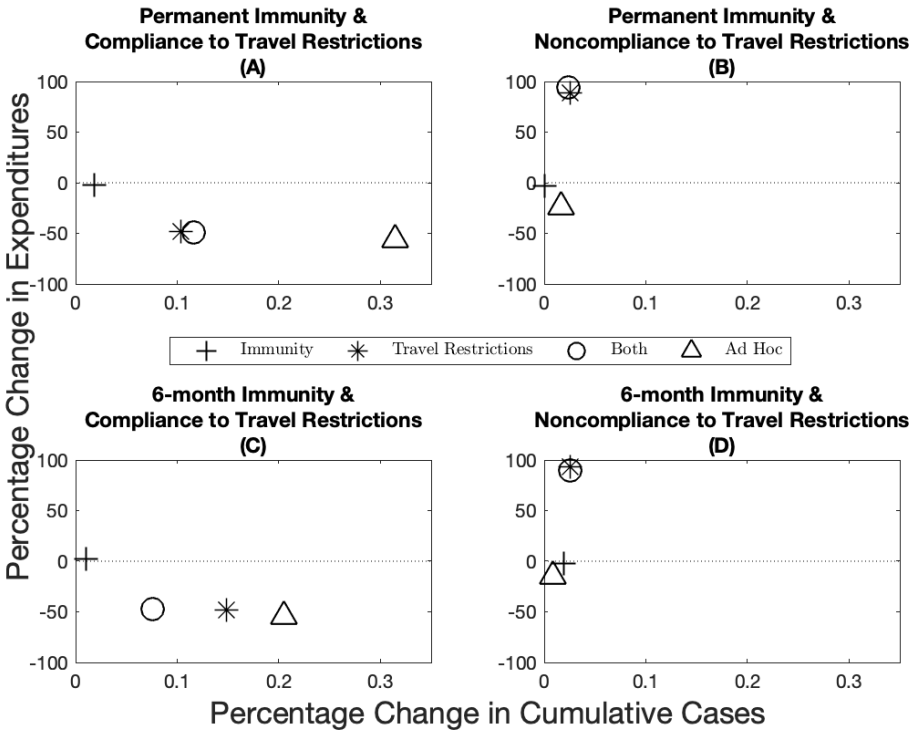

Figure 9: **Robustness of epidemiological and economic outcomes under different scenarios with vaccines.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 10%. The x-axis represent small percentage changes but when scaled up to population level effects translate into significant differences in public health outcomes.

We also investigate the robustness of the optimal allocations when the demographic characteristics are heterogeneous across jurisdictions. When jurisdictions have a different case-fatality ratio, the *ad hoc* allocation performs better than the optimal allocations when considering cases as the main health outcome (Figure B.15). However, this approach is misleading because when case-fatality ratios are heterogeneous across jurisdictions, the cumulative aggregate number of cases (all jurisdictions together) is a poor outcome measure as a case in one place is not the same as a case in another jurisdiction. In this setting, the disease burden and cumulative damages give a more accurate depiction of the situation. In fact, while the *ad hoc* allocation outperforms the optimal allocations in terms of cumulative cases, it performs considerably more poorly when considering cumulative damages. We generally find that the optimal allocations outperform the *ad hoc* allocation

in all scenarios considered (Figure B.16). When jurisdictions have a different contact rate, the same pattern as in Figure 9 holds in the sense that when there is compliance to travel restrictions, the optimal allocations outperform the *ad hoc* allocation, while the *ad hoc* generally performs better than the optimal allocations when there is noncompliance to travel restrictions (Figure B.17).

2.3.4 Sensitivity Analyses

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

Two key parameters in our analysis are the scale of the workability cost (c_A in Equation (9)) and the level of vaccine effectiveness (see Appendix B.3 for more details). While imposing the *ad hoc* rule *ex ante* implicitly means that the cost of deviating from the *ad hoc* allocation is infinite, in practice it is likely finite but hard to quantify, as it depends on logistical, political, and cultural factors. We investigate the sensitivity of our results by solving for optimal vaccine allocation over a range of values. We find greater deviations off of the *ad hoc* at lower workability costs resulting in greater differences in cumulative cases, and smaller deviations as the workability cost parameter increases (Figure B.18 panels A, B, C, and D). Specifically, we find that when the cost is in the neighborhood of the VSL (c in Equation (7) and Figure B.18 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 [101] (see Appendix B.1 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 B.19 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 B.19 panels A, B, C, and D).

2.4 Conclusion

Recent studies have discussed how a vaccine against the coronavirus disease (COVID-19) should be allocated within a geographical area (see for instance [58; 59; 60]) and on a global scale (see for instance [61; 62; 63; 64]). Building off the spatial-dynamic literature in epidemiology, we contribute to this body of work by addressing the question of distributing a relatively scarce COVID-19 vaccine across smaller geographic areas, such as counties, regions, or states. The U.S. National Academies of Sciences, Engineering, and Medicine (NASEM) [65] recommends to allocate a 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 such *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, the vaccine capacity constraint, and the demographics across jurisdictions are varied. In most of these scenarios, we find that priority should be given to jurisdictions that initially have lower disease burden. The intuition behind this result—already put forward by Rowthorn et al. [47] 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 [102]. We find that higher vaccine capacity can lead to the opposite result for a short period of time at the beginning of the time horizon, and we find that the high burden jurisdiction should be prioritized when it has more essential workers, as long as its infection level is increasing and remains higher than the jurisdiction with fewer essential workers. Our results also suggest that deviations from an *ad hoc* allocation rule based on population size are highly beneficial when one jurisdiction has a population with higher mortality rates.

We also show the value of complying to a travel restriction, as compliance leads to lower cumulative damages across both jurisdictions, regardless of whether immunity is permanent or temporary. The reduction in cumulative damages is particularly important for the jurisdiction with fewer infected individuals. Considering nonlinear damages due to an overload of health care systems [103; 104] and a corresponding varying death rate due to scarce intensive care unit beds [75], and other second-order problems such as consumption losses [105; 106; 107], excess mortality [108], and psychological distress [109] 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 for a wide range of values we considered; this result holds in all scenarios we considered in our analysis. We considered *ad hoc* allocation rules that favor "speed and workability" (put forward by NASEM [65]). Other allocation rules are possible. For instance, in the base case of our paper, we assumed identical contact rates across jurisdictions. In turn, this implied that the movement within a given jurisdiction is assumed to be identical across jurisdictions. In practice, population mobility likely differs from one jurisdiction to another and an *ad hoc* allocation could be based on population mobility and contact structure. The methodology employed in this paper can investigate the trade-offs of other *ad hoc* rules and as a result, can offer potentially important information to policymakers that face the challenge of allocating scarce COVID-19 vaccines to their jurisdictions.

Extrapolating our results to the entire U.S. suggests that allocating a vaccine based on the *ad hoc* allocation rule advocated for by NASEM can have serious public health consequences. How many additional cases accrue 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 vaccination 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.

There are, however, important factors that have received significant attention in the literature that we should fully incorporate in future research. For example, the composition of the population of a jurisdiction is assumed to be homogeneous. To mimic the fact that the virus disproportionately affects elderly people [67] and/or people with pre-existing conditions [110], and also to mimic the fact that the risk of infection is highly occupation dependent [68], we simply assumed that the average case-fatality ratio and contact rate was higher in one jurisdiction. In practice, however, the composition of a population within a given jurisdiction is not homogeneous. Further research combining heterogeneity both across jurisdictions in the form of different disease burden and within jurisdictions in the form of different risk of complications and risk of infection could add additional valuable insights into the trade-offs inherent in these different allocations rules.

Finally, while our paper and most of the discussion revolves around the allocation of a vaccine, a similar allocation problem may arise if an antiviral drug were to become available (for a discussion on antiviral treatments for SARS-Cov-2, see [111]). Because

drugs and vaccines have different goals—treating infected individuals and prophylaxis, respectively—the economic and public health trade-offs of different allocation rules may be unique to the type of pharmaceutical intervention. Future work considering the joint allocation question of antiviral drugs and vaccines could be valuable in understanding the trade-offs and complementarities between these different pharmaceutical interventions.

3 Chapter 3: "Ecological Interventions to Fight Diseases: When Objectives are Misaligned"

Abstract

Ecological interventions are nonpharmaceutical actions that individuals and communities can take to reduce disease transmission by considering interactions among the different disease-carrying organisms and their environment. They are receiving increasing interest by both researchers and practitioners to prevent and manage environmentally transmitted diseases because of their potential to maintain control of the disease burden in a decentralized way. When the main drivers of such interventions are the economic benefits derived from undertaking the activity, ecological interventions also have the potential to be sustainable. While economic incentives are key in making such interventions viable, they can also have unintended consequences that hinder their public health purpose. In this paper, we investigate the consequences that profit-maximizing behavior has on the transmission of an environmentally transmitted disease using schistosomiasis and prawn aquaculture as a case study. Freshwater prawn, a natural predator of the intermediate host snail, is a food source for local human populations, and it maintains disease prevalence at low levels in the environment. We simulate the behavior of a profit-maximizing aquaculture farmer whose activity has the potential to create ancillary public health benefits, but whose decision-making concerning supplemental feeding can mitigate these same public health benefits. Our results suggest that feeding prawns to maximize profits can lead to a divergence between the health- and profit-maximizing outcomes. We investigate the performance of two policies (standardized rotation length and a limited feeding season) to reduce this divergence and realize the public health benefits from prawn aquaculture.

Key words: economic epidemiology, applied numerical methods, optimal timing, aquaculture regulation.

3.1 Introduction

Vaccines, drugs, and pesticides are conventional tools used to fight infectious diseases. These interventions target the risk in humans (e.g. vaccines, drugs) or their immediate environments (e.g. pesticides) [56]. Despite uncontestable public health and economic benefits attributable to these pharmaceutical interventions for some diseases (for instance, see [112] for smallpox and [113; 114] for Covid-19), they show limited effectiveness against other diseases that are environmentally transmitted and have complex life cycles involving non-human host [5].

An example of a disease with such a complex life cycle is schistosomiasis, a debilitating parasitic disease that affects over 250 million people worldwide [115], and its transmission relies on a non-human intermediate host (a freshwater snail) [35]. The parasite's life cycle requires a phase of asexual replication within the intermediate host, resulting in the shedding of cercariae (the larval form of the parasite) that contaminates rivers, lakes, or other bodies of freshwater with which humans may be in contact [35]. Conventional interventions to fight schistosomiasis—including a drug that effectively eliminates the parasites in humans [6] and molluscicides that eliminate the intermediate host in the environment [54]—have their share of problems. Despite being developed in the mid 1970s, investments in the billions of US dollars in mass drug administrations have had little effect on the global disease burden [5], likely because treated individuals often have no other option but to continue using to parasite-contaminated waters [4]. The molluscicides used in bodies of freshwater where snails are present are also toxic to other species [35; 48; 49; 50] and could be prone to resistance evolution [56] (for an example of resistance evolution for insecticides used for malaria control, see [55]). Combining multiple types of treatment can reduce the overall cost of treatment and the disease burden [5; 12; 19; 116], but separate guidelines for each type of treatment, such as current World Health Organization (WHO) recommendations, considerably increase the time span over which drug treatment needs to be administered [100].

Because of these numerous problems, researchers and practitioners are considering ecological approaches to address these challenges. Ecological interventions are based on the complex interactions between human and non-human hosts, and how they interact with their environments [56]. Examples include fertility control of the intermediate host using, for example, gene-drive technologies (see [117] for the use of such technologies to fight malaria and see [118] for the potential of such technologies to fight schistosomiasis), natural habitat manipulation (see [119; 120; 121; 122] for examples relating to schistosomiasis), or the use of predators of the intermediate host as biological controls (the focus of this paper). Introducing non-native predators reduces schistosomiasis prevalence in humans [123], but using non-native species as biological controls can be risky due to potential catastrophic nontarget effects [124]. Modern biological controls can be safe and effective when done properly and when using native species [125]; an example of this in the context of schistosomiasis is the restoration of native freshwater prawns (a natural predator of snails), which reduces snail-to-man transmission in areas where schistosomiasis is endemic [116]. A recent study [42] showed that prawn aquaculture has the potential to reduce the snail population, aid in schistosomiasis control, and be a source of sustainable development.

An important advantage of prawn aquaculture to fight schistosomiasis is that the economic incentives can be the driving force of the intervention. The public health benefits that prawn aquaculture provides could be a positive externality created by the economic activity. To sustain decentralized management of this ecological approach to fight schistosomiasis, economic incentives are key; otherwise there is no guarantee of having participation of local communities in the endemic areas. While prawn aquaculture as a means to fight schistosomiasis may seem like a win-win situation if the economic incentives are present, issues may arise if the public health objective differs from the private profit-maximizing objective (i.e., if there is a misalignment of objectives between the central planner and the decentralized economic agents). The reason why a divergence may occur is that the ones who bear the costs of aquaculture do not incur all the benefits of the activity. For instance, the aquaculture farmers are not the only ones benefiting from the

reduction in disease transmission. The aquaculture farmers may also not consider (or realize) the health benefits that they themselves reap from their own activities.

In this paper, we explore a potential mechanism that could lead to a divergence between the health-maximizing optimum and the profit-maximizing optimum. Specifically, we simulate a situation where the aquaculture farmer can supplement the prawns' diet with feed. The supplemental feed increases prawn growth but reduces predation of snails, which in turn can mitigate the public health benefits associated with prawn aquaculture. We use optimal control techniques to compare the dynamics of infected intermediate hosts and infected humans in scenarios that differ because of the objective of the decision-maker.

Our findings suggest that ignoring potential investments in optimizing the returns from prawn aquaculture due to a profit-maximizing behavior (i.e. ignoring the fact that farmers may want to feed prawns to increase their profits) could lead to an important divergence between the public health and the private optima. The extent to which these two outcomes differ inherently depends on (i) the degree to which prawns substitute away from snails when they are fed, and (ii) the extent to which prawns are generalist predators. We investigate how these different aspects affect epidemiological outcomes, and we present different policies that could be implemented to incentivize aquaculture farmers to make decisions that result in a greater reduction in disease transmission.

This paper is divided as follows. In Section 3.2, we detail the components of the predator–prey model, the epidemiological model, and the economic objective of the decision-maker. Section 3.3 presents the base case results, the simulated policies meant to reduce the impact of feed, and the sensitivity analyses of key parameters. Section 3.4 concludes the paper.

3.2 Model

To illustrate the challenges of using prawn aquaculture for public health goals, we develop a stylistic model for a single representative prawn aquaculture enclosure. A nearby small village (5,000 people, following Hoover *et al.* [42]) is assumed to come into contact with

freshwater only within the boundaries of the aquaculture enclosure. The snail population is assumed to be closed within the enclosure (i.e. there is no movement in or out of it) and the larval stage of the parasite living in freshwater is assumed stay within the boundaries of the enclosure and not be able to migrate into it. The objective consists of optimizing the harvest time of the prawns, where their growth is affected by both exogenous factors (i.e. the predation of snail and non-snail prey) and endogenous factors (i.e. supplemental feed).

There are many analogies between the profit-maximizing aquaculture problem studied in this paper and optimal rotation lengths in forestry. The trade-off faced by the decision-maker is between the marginal value of waiting an extra period of time (in our case, the additional aquaculture profits earned by allowing prawns to grow one more day) and the opportunity cost of the aquaculture operation (i.e. the foregone benefits of not harvesting today); the optimal rotation length is given by the point in time when these two objects are equal to each other. When farmers do not feed the prawns, this trade-off is solely affected by biological factors (i.e. the prawns' growth rate) and time preference (i.e. the discount rate). When farmers feed prawns (analogous to silviculture in forestry), feed affects both sides of the equation. Everything else equal, feed increases the marginal value of waiting an extra period of time since the biomass (and thus profits) tomorrow will be larger if prawns are fed, and it also increases the opportunity cost of the aquaculture operation since supplemental feeding leads to higher intertemporal profits. The last—but yet very important—analogy to be made with forestry is that prawns procure ancillary public health benefits, analogous to amenity values in forestry problems; see e.g. [126]. By being a generalist predator [127], prawns consume the freshwater snails, which spread the pathogens in the freshwater where villagers go for their everyday activities. This may create a divergence between the health- and profit-maximizing outcome if the private decision-maker does not take this into account.

We solve the problem for different objectives. Specifically, we consider (i) a private objective where only aquaculture profits are maximized and (ii) a public health objective where avoided infection costs are maximized net of stocking and feeding costs.

3.2.1 Predator–Prey Model

The predator–prey model of this paper is inspired by the work of Hoover *et al.* [42]. The main difference is that we explicitly model two components of the prawns' diet: freshwater snails (N) and supplemental feed (U). The prawns' somatic growth is

$$\dot{L}(t) = \left(\frac{\max\left\{k, \alpha_N(t)N(t)\right\} + \alpha_U\sqrt{U(t)}}{1 + g\Omega(t)} \right) \left(L_\infty - L(t) \right) \tag{11}$$

where $L(t)$ is the average length of prawns at time t, L_∞ is the asymptotic length of prawns, k is the intrinsic growth parameter, α_i with $i = N, U$ is the attack rate of prawns on snails ($i = N$) and supplemental feed ($i = U$), and g is parameterized to mimic a reduction in somatic growth at higher biomass $\Omega(t)$ levels [128]. Note that while the prawns' attack rate on snails $\alpha_N(t)$ varies over time due to changes in prawn-to-snail body size ratio [42; 129], the prawns' attack rate on supplemental feed α_U is assumed constant. The number of prawns $P(t)$ varies according to some mortality rate that exponentially decreases with increasing body size [130] and that is dependent on biomass because of a density-dependent competition for food [128]. The prawn enclosure is assumed to contain all males [131; 132] or all females [133], meaning that prawns do not reproduce. See Equation (C.3) in Appendix C.1 for more details on the dynamics of the prawn population.

The prawns' predation is dependent on the relative abundance of a prey, its handling time, and the attack rate. We illustrate the case for the $W(t)$ infected snails below, but the per-capita predation rate for the $X(t)$ susceptible snails can be found by substituting the numerator accordingly; note that the sum of the infected and susceptible snails is equal to the total population of snails, i.e. $W(t) + X(t) = N(t)$. The per-capita prawn predation rate of snails is modelled as a Holling type III functional response [134]:

$$\psi_W(t) = \left(\frac{\alpha_N(t)\epsilon W(t)^n}{1 + \alpha_N(t)T_{hN}(t)\epsilon N(t)^n + \alpha_U T_{hU} U(t)^n} \right) \tag{12}$$

where T_{hi} for $i = N, U$ is the prawns handling time of snails ($i = N$) and supplemental feed ($i = U$), ϵ represents additional searching costs not present in laboratory settings where parameters were estimated [129], and n represents the exponent of Holling's type III functional response. Note that a Holling type III functional response implicitly accounts for other non-snail prey in the prawns' diet [134]; see [42; 135] for other applications of a Holling type III functional response to prawn predation of snails. Like their respective attack rates, the prawns' handling time of snails $T_{hN}(t)$ varies over time due to changes in prawn-to-snail body size ratio [42; 129], but the prawns' handling time of supplemental feed T_{hU} is assumed to be constant.

3.2.2 Model of Disease Transmission

The dynamics of schistosomiasis are modeled using a SIS-W model, which tracks the dynamics of the fraction of susceptible humans (S), the fraction of infected humans (I), and the dynamics of the infected intermediate hosts (W) in the environmental reservoir (i.e. the aquaculture enclosure next to a village where the disease is endemic); see [4] for a general model of environmental disease transmission. Susceptible humans, which we denote as $(1 - I)$ to reduce the number of variables, contract the parasite through contact with the contaminated reservoir at a rate of β; the fraction of infected humans goes down over time from natural recovery [27]. The fraction of infected humans grows according to:

$$\dot{I}(t) = \beta W(t)\Big(1 - I(t)\Big) - \gamma I(t) \tag{13}$$

where γ is the natural recovery rate.

Infected humans shed parasites in the environment at rate of λ, and the shed parasites in turn infect the susceptible snails. The snails grow according to a logistic-type reproduction function, and only healthy snails can reproduce since the parasite castrates the snails [12; 42; 100; 116]. The infected snails are either consumed by the prawns or die at a disease-induced mortality rate of δ. Note that intrinsic density-dependent mortal-

ity is implicitly accounted for in the nonlinear term of the logistic growth equation; see Equation (C.4) in Appendix C.1 for more details. The growth of infected snails is:

$$\dot{W}(t) = \lambda I(t)X(t) - \delta W(t) - \psi_W(t)P(t) \tag{14}$$

where $\psi_W(t)P(t)$ represents the total predation of infected snails.

3.2.3 Profit- and Health-Maximizing Objectives

The aquaculture farmer's objective is to maximize profits. We assume the farmer only considers their own profits and does not take into account the indirect public health benefits of aquaculture (i.e. the reduced disease transmission). The farmer must incur a startup cost to purchase the juvenile prawns (at a per-unit cost of c_P) that will grow in the enclosure; here we must assume the farmer has the means to purchase the juvenile prawns. To maximize profits, the farmer choose the optimal date T^* on which to harvest the prawns, and potentially, to determine the prawns' optimal feeding path over the endogenously determined time horizon (Appendix C.1 details the cost function defined as $c_U\big(U(t)\big)$). After growing for T^* days, the prawn biomass $\Omega(T^*)$ is sold at a price of p per kg; revenues and feeding costs are discounted at a nonnegative rate of r. The single-rotation profit-maximizing objective is

$$\max_{T,U(t)} e^{-rT}p\Omega(T) - c_P P(0) - \int_0^T e^{-rt}c_U\big(U(t)\big)\ \mathrm{d}t \tag{15}$$

subject to the predators' dynamics (equations 11 and C.3), the preys' dynamics (Equation C.4), and the dynamics of disease transmission (equations 13 and 14), along with non-negativity and physical constraints on the predators, preys, and feed, and the free endpoint conditions (see Appendix C.2 for more details).

The single-rotation objective described above is a simplified version of the actual problem. Using this objective essentially assumes that the decision-maker chooses the rotation length and feeding path without considering that another rotation will start anew once the current one finishes. Given the short length of the aquaculture rotations (see e.g. [42],

where optimal rotation lengths vary between 164 and 331 days when prawns are not fed), we need to consider a multi-rotation objective. When pushed to the limit—that is, when the farmer cares about an infinite number of rotations—the objective function needs to be modified accordingly. This is the profit-maximizing objective of interest in the paper; see Equation (C.8) in Appendix C.2. Note that we assume constant parameters over time, so they do not change from one rotation to another, which implies every rotation is identical.

The health objective maximizes the avoided infection costs, net of any feeding and stocking costs. The avoided infection costs are defined as being the difference in discounted damages between the beginning and ending infection levels. For illustration purposes, we present the single-rotation health-maximizing objective; it is given by:

$$\max_{T,U(t)} \; c_I\Big[I(0) - e^{-rT}I(T)\Big] - c_P P(0) - \int_0^T e^{-rt} c_U\big(U(t)\big)\; \mathrm{d}t \tag{16}$$

where c_I is a cost parameter representing the loss of disability-adjusted life years (see [12] for an estimate of the number of disability-adjusted life years lost). The health objective is subject to the same constraints as the private objective defined above. As in the profit-maximizing case, our objective of interest in the health-maximizing case is the infinite-horizon harvest, which is given by Equation (C.9) in Appendix C.2.

3.3 Results

We solve for the optimal rotation length and the optimal feeding path over the endogenously determined time horizon. In particular, we focus on two different scenarios: (i) the infinite-horizon health-maximizing objective and (ii) the infinite-horizon profit-maximizing objective. We investigate the effect of the different objectives by mapping out the dynamics of the aquaculture profits, the infected snails, and the feeding paths.

We use pseudospectral collocation to solve for the optimal harvest time and optimal feeding path. This method for solving optimal control problems converts the continuous time problem into a constrained non-linear programming problem and finds the coeffi-

cients that best fit the polynomial at the endogenously determined collocation nodes; see [45; 100; 136] for other applications of this technique and see Appendix C.2 for more details.

3.3.1 Base Case: No Central Planner Intervention

The results using our preferred specification of the parameters (see Appendix C.1 for more details) are presented in Figure 10. Panel A shows the aquaculture profits of one rotation as a function of time (in days) for the health-maximizing objective (in blue) and for the profit-maximizing objective (in red). Panel B shows the dynamics of the infected freshwater snails (as proportion of their steady-state) as a function of time (also in days) for the same two objectives. Panel C shows the health- and profit-maximizing feeding paths. In each of the three panels of Figure 10, the blue and red lines stop at their corresponding optimal objective-dependent rotation length.

From Figure 10 we note that farmers have a private incentive to feed prawns (Figure 10, Panel A; the red line, here representing profits, is higher than the blue line) because it increases the growth rate of prawns (see Equation 11) and in turn leads to higher profits. There is a social cost associated with feeding that they do not account for (Figure 10, Panel B; the red line, here representing infected snails, is higher than the blue line) that stems from the fact that feeding prawns decreases prawn predation on snails (see Equation 12).

We see two main forces at play when looking at the feeding paths (Figure 10, Panel C). On the one hand, health-related benefits (i.e. the avoided infection costs) pushes the health-maximizing decision-maker to feed prawns relatively more at the beginning and ending of the time horizon (i.e. the health-maximizing feeding path is U-shaped; see the blue line in Figure 10, Panel C). On the other hand, aquaculture benefits pushes a profit-maximizing decision-maker to gradually increase feed at the beginning of the time horizon until reaching a peak and then slowly reducing until the optimal stoppage time (i.e. the profit-maximizing feeding path has an inverse U-shape; see the red line in Figure 10, Panel C). Note that while the pattern differs between the health- and profit-maximizing

feeding paths, the magnitude of feed remains considerably higher in the profit-maximizing case (in Figure 10, Panel C, there is a separate y-axis for each objective).

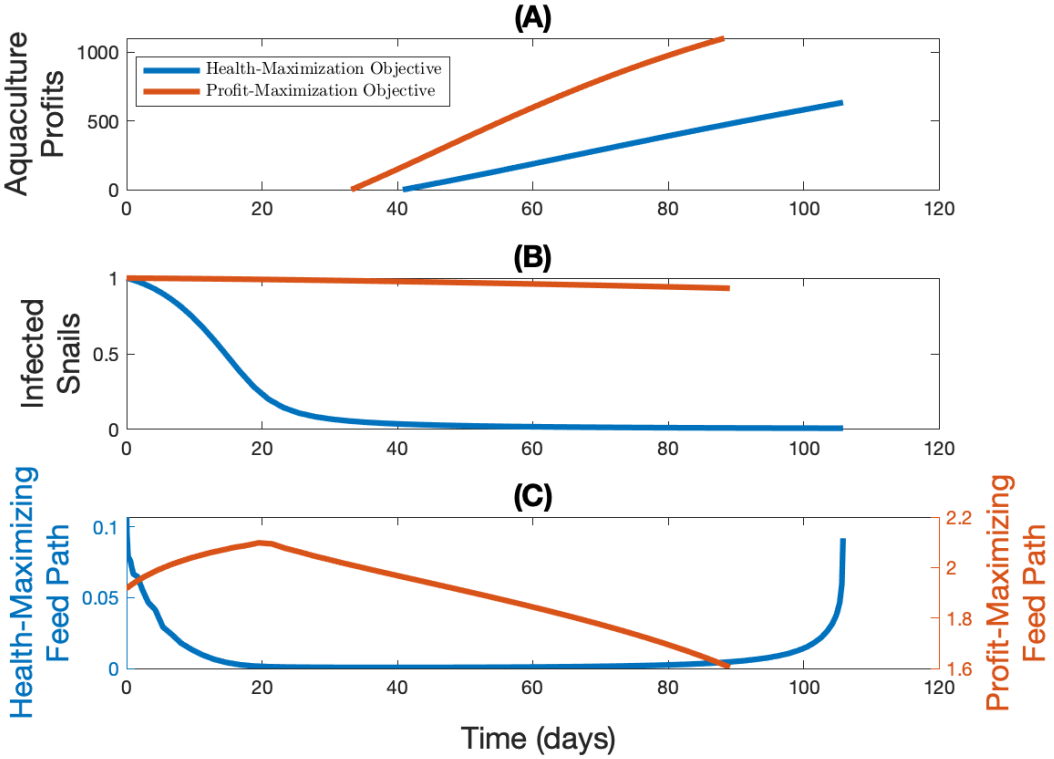

Figure 10: **Economic and epidemiological outcomes under the health-maximizing (in blue) and profit-maximizing (in red) objectives.** This figure shows the dynamics (in days) of the discounted aquaculture profits of one rotation (Panel A), the corresponding number of infected snails as a proportion of its initial steady-state value (Panels B), and the health-maximizing feeding path (Panel C, main y-axis in blue) and profit-maximizing feeding path (Panel C, secondary y-axis in red).

These results indicate that the objective of the decision-maker can greatly affect epidemiological outcomes. Because of supplemental feed, a profit-maximizing decision-maker could create an important divergence between the health- and profit-maximizing outcomes that the previous literature did not identify [42]. This shows the importance of better understanding the implications of aquaculture profit-maximization.

Table 3 summarizes the impact of the divergence between the health-maximizing outcome and the profit-maximizing outcome. It shows (i) the different rotation lengths (in days), (ii) the average snail density at the end of the rotation, (iii) the additional human cases per 1M people attributable to profit-maximization, and (iv) the aquaculture profits of a

single rotation. We can see that the health- and profit-maximizing objectives lead to an important divergence in average final snail densities, which results in important human health impacts. Because of this important divergence, it takes approximately 5.65 years for the profit-maximizing snail density to reach similar levels as the health-maximizing snail density, which results in 7,760 additional human cases per 1 million people. Finally, fed prawns grow faster and provide considerably more profits than non-fed prawns, so a profit-maximizing decision-maker will be incentivized to feed.

Objective	**Rotation Length (Days)**	**Average Final Snail Density (per m^2)**	**Additional Human Cases (per 1M People)**	**Aquaculture Profits (USD)**
Health	106	0.48	–	633
Profits	89	46.98	7,760	1,112

Table 3: **Impact of the divergence of outcomes between the health- and profit-maximizing optima.** This table shows the rotation length (in days), average final snail density (per m^2), additional human cases due to profit-maximization (per 1M people), and aquaculture profits (USD) for the health- and profit-maximizing outcomes with and without supplemental feed.

3.3.2 Social Planner Intervention and Sensitivity Analyses

In this section, we simulate two different policies that a central planner could implement to limit the impact of profit-maximization on health outcomes. Specifically, we impose (i) a minimum, or standardized, rotation length and (ii) a limited feeding season. By imposing these policies and by varying parameters that play a key role in determining the rotation length, we investigate how these policies affect the dynamics of the key economic and epidemiological outcomes—the aquaculture profits, the population of infected snails, and the feeding paths.

In both the minimum rotation length and limited feeding season, the regulation is based on the optimal rotation length of the health-maximizing objective. In the minimum rotation length policy, the profit-maximizing rotation length needs to be at least as long as the health-maximizing one. Hence, we first find the health-maximizing rotation length and we impose it as a constraint to solve the profit-maximizing case. Because feeding reduces prawn predation of snails, this regulation does not guarantee that the snail pop-

ulation in the health- and profit-maximizing cases are identical across the no-policy and policy cases, but it does guarantee that the shortness of the profit-maximization rotation is not the driving force of the divergence.

For the limited feeding season, we model it by banning feeding for the latter part of the rotation. As before, we first find the health-maximizing rotation length, but this time, we allow feeding to occur only for the first half of the health-maximizing rotation length. Compared to the minimum rotation length regulation, there is no constraint on how long the rotation needs to be so the divergence between the health- and profit-maximizing cases can be due to shorter rotations and reduced predation of prawns on snails.

Efficiency of Prawns at Converting Supplemental Feed

Figure 11 shows the effect that varying prawn feed conversion efficiency—i.e. the effect feed has on prawn growth— has on the aquaculture profits of a single rotation. It presents different cases where there is no intervention from the social planner (Figure 11, panels A, D, and G), where the social planner imposes a standardized rotation length (Figure 11, panels B, E, and H), and where the social planner limits the feeding season (Figure 11, panels C, F, and I). We see that rotation lengths and the aquaculture profits of a single rotation increase as the prawns' feed conversion efficiency increases (Figure 11; compare the height and length of the upper panels with the height and length of the lower panels). While health-maximizing rotation lengths are shorter than profit-maximizing rotation lengths at lower levels of feed conversion efficiency (e.g. compare rotation lengths in Panel A of Figure 11), the opposite is true for high levels of feed conversion efficiency (e.g. compare rotation lengths in Panel G of Figure 11).

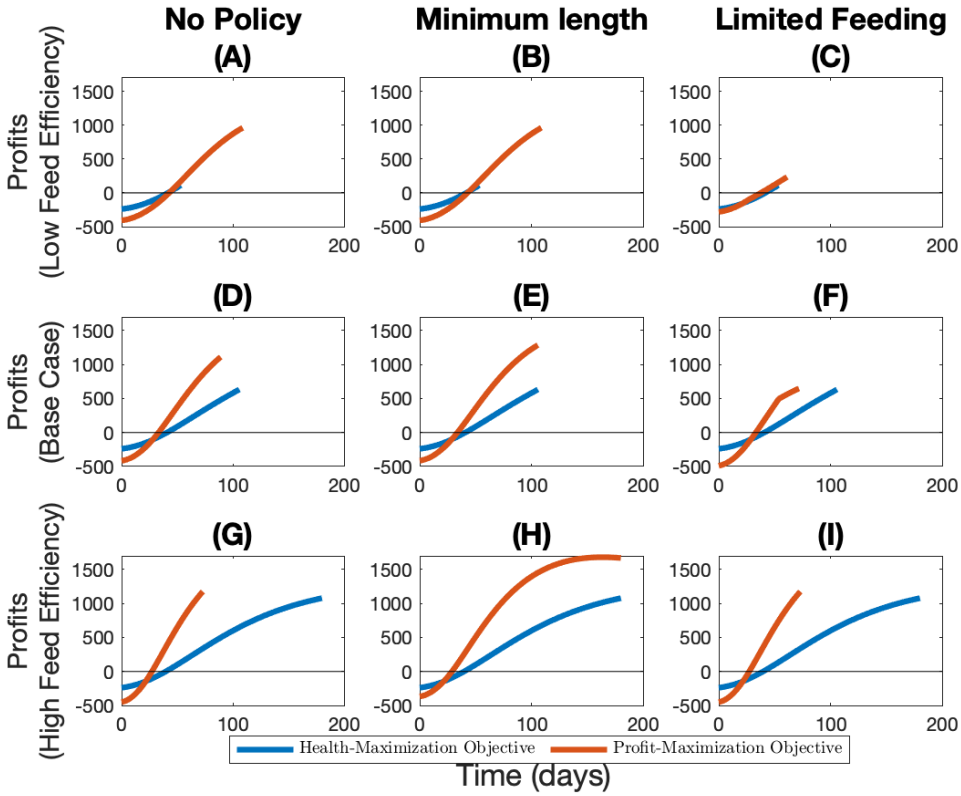

Figure 11: **Dynamics of the aquaculture profit functions under different policies and for different feed conversion efficiencies.** The figure depicts the change in harvesting profits when there is no policy (panels A, D, and G), with a minimum rotation length policy (panels B, E, and H), and with a limited feeding season policy (panels C, F, and I) for the health- (in blue) and profit-maximizing (in red) objectives. Panels A, B, and C represent the case where prawns have a lower feed conversion efficiency than in the base case (i.e. 50% lower); panels D, E, and F represent the base case feed conversion efficiency; panels G, H, and I represent a case where prawns have a higher feed conversion efficiency than in the base case (i.e. 50% higher). See Figure C.1 for an alternative representation of these results.

Figure 12 presents the effect on snail dynamics of varying the prawns' efficiency to convert supplemental feed. When there is no policy in place, higher feed conversion efficiency implies a lower predation of snails, which implies worsened health outcomes (the red line gets less and less steep as we go from Panel A to Panel D, and from Panel D to Panel G) because more and more feed is used (red lines with secondary y-axis; Figure 13, panels A, D, and G). The minimum rotation length policy is only effective at higher levels of feed conversion efficiency (Figure 12, panel H); it has very little, or no, effect at lower levels when health- and profit-maximizing rotation lengths are similar. Compared to the minimum rotation length, the limited feeding season performs better at lower levels of

feed conversion efficiency because the health-maximizing rotation length is short and so feeding is only allowed for a very short period of time (see e.g. Figure 12, Panel C). When feed conversion efficiency is higher, the health-maximizing rotation length is much longer and more than double that of the profit-maximizing rotation, meaning that the policy has no effect. The limited feeding season policy generally seems to yield a more health-maximizing result than the minimum rotation length (i.e. the red line ends up being closer to the blue one in the third column (Figure 11 panels C, F, and I) than in the second column (Figure 11 panels B, E, and H)). The reason for this is that if the marginal value of waiting an extra period is still higher than the foregone benefits of not harvesting today when feeding is banned, the predation of prawns on snails will be very high (as prawns will be considerably larger than their juvenile size) and the snail population will be eliminated very quickly (e.g., see Figure 12, panels C and F).

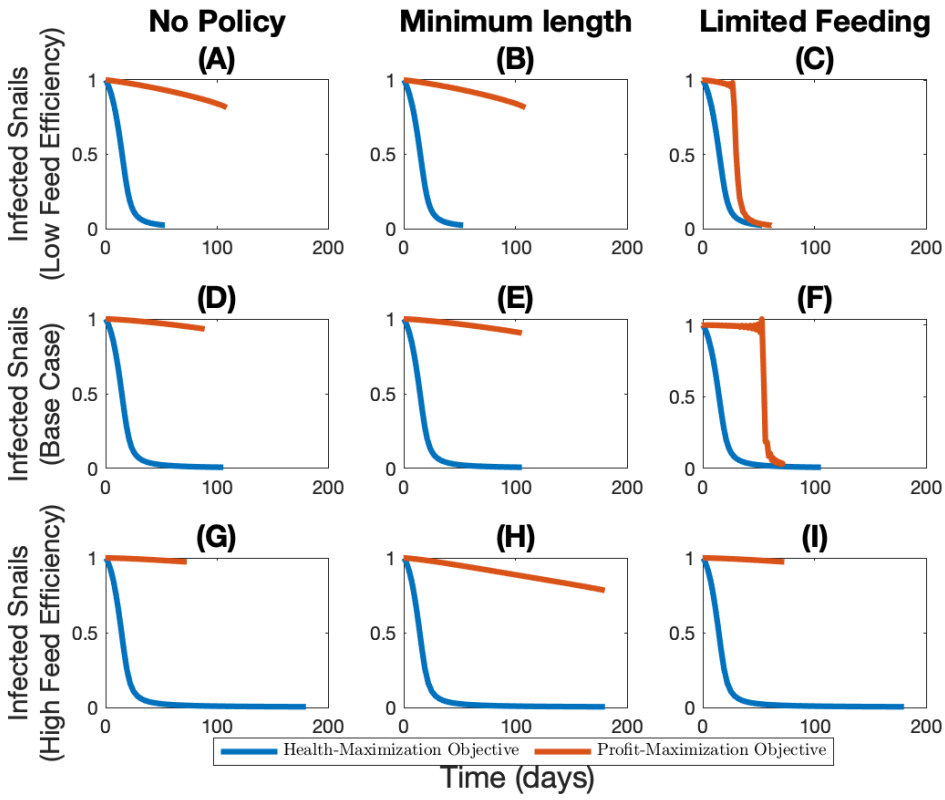

Figure 12: **Dynamics of the infected snail population under different policies and for different feed conversion efficiencies.** The figure depicts the change over time in the number of infected snails as a proportion of its steady-state value when there is no policy (panels A, D, and G), with a minimum rotation length policy (panels B, E, and H), and with a limited feeding season policy (panels C, F, and I) for the health- (in blue) and profit-maximizing (in red) objectives. Panels A, B, and C represent the case where prawns have a lower feed conversion efficiency than in the base case (i.e. 50% lower); panels D, E, and F represent the base case feed conversion efficiency; panels G, H, and I represent a case where prawns have a higher feed conversion efficiency than in the base case (i.e. 50% higher). See Figure C.2 for an alternative representation of these results.

Figure 13 shows the feeding paths for different feed efficiency conversions and for the different policies. Without any decision-constraining policies, a higher feed conversion efficiency implies a higher level of feed for both the health- and profit-maximizing objectives (Figure 13, panels A, D, and G). However, the policies affect the feeding path differently. In anticipation of the limited feeding season, profit-maximizing decision-makers increase feed initially to compensate (red lines with secondary y-axis; Figure 13 panels C and F). When the minimum rotation length policy is most effective (i.e. in Figure 13, Panel H), the feeding path is reduced compare to a case where no policy is in place (Figure 13,

Panel G) because the farmer knows it will have to wait until harvesting and does not want to reach the maximum of the profit function as quickly as it would have without any regulation.

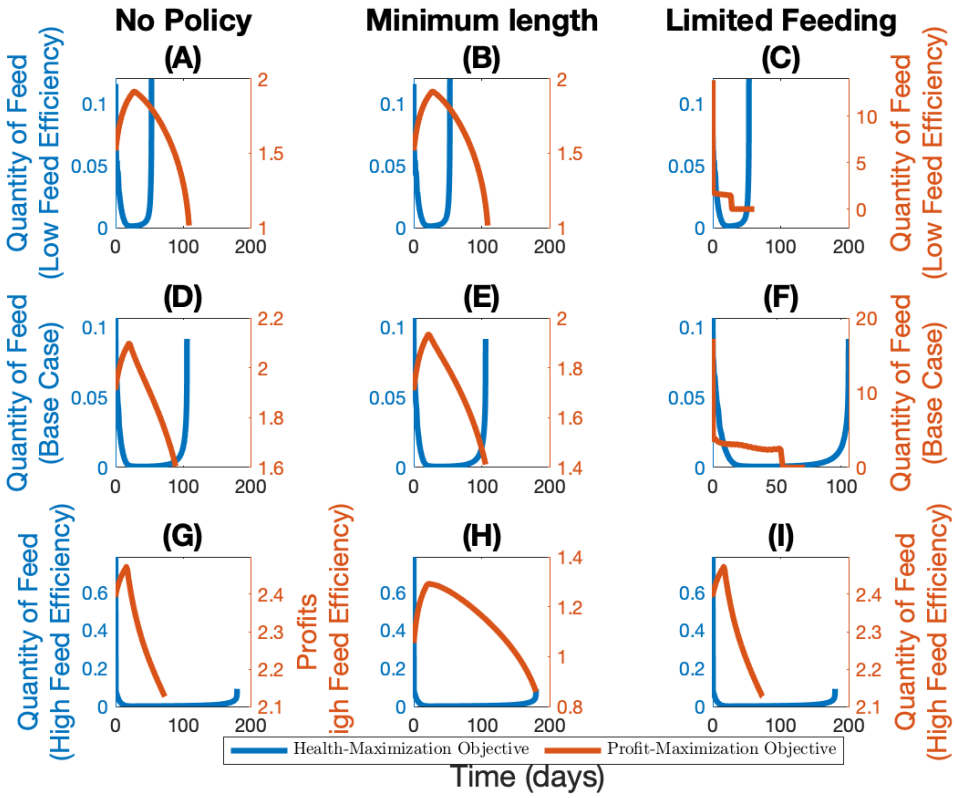

Figure 13: **Dynamics of the feeding paths under different policies and for different feed conversion efficiencies.** The figure depicts the change over time in feed when there is no policy (panels A, D, and G), with a minimum rotation length policy (panels B, E, and H), and with a limited feeding season policy (panels C, F, and I) for the health- (in blue) and profit-maximizing (in red) objective. Panels A, B, and C represent the case where prawns have a lower feed conversion efficiency than in the base case (i.e. 50% lower); panels D, E, and F represent the base case feed conversion efficiency; panels G, H, and I represent a case where prawns have a higher feed conversion efficiency than in the base case (i.e. 50% higher). See Figure C.3 for an alternative representation of these results.

Specialist or Generalist Predators

Figures C.4, C.5, and C.6 in Appendix B.3 show respectively the dynamics of the aquaculture profits, the infected snail population, and the feeding path when the exponent of the Holling type III functional response is varied. A smaller exponent mimics a situation where prawns are more specialist predators (i.e. rely more on feed and snails) and a higher exponent mimics a case where prawns are more generalized predators (i.e. rely less

on feed and snails) [134]. The more generalist predators that prawns are, the lesser the impact of feed, which in turn implies a longer health-maximizing rotation (Figure C.4; compare the rotation length between e.g. Panel A and Panel C) and a higher feeding path (Figure C.6; compare the blue feeding paths of e.g. Panel A and Panel C). Since feed conversion efficiency is held constant throughout (at our base case), the degree to which prawns are specialist or generalist predators has practically no impact when prawns are fed to maximize profits (Figure C.4; compare the rotation length and red profit functions between panels A, B, and C).

Generally, the more generalist predators that prawns are, the less snails play a role in their diet, and the higher the final snail population will be (Figure C.5; compare the red lines of e.g. Panel A and Panel C). Because the more generalist prawns lead to longer health-maximizing rotations, the minimum rotation length policy is most effective when prawns are more generalist predators (Figure C.4; e.g. the red lines in panels A and B and those in D and E are very similar). The limited feeding season policy performs poorly when prawns are more specialist predators, because the policy is not binding. The profits are considerably lower than in the no-policy case (Figure C.4; compare panels A and C), and the policy does not have the anticipated effect in terms of reducing the population of infected snails (Figure C.5, Panel C). The health-maximizing rotation length, on which the policy is based, is considerably shorter, which pushes a profit-maximizing agent to feed a very large amount of feed initially and then harvest at the point in time when the ban on feed starts.

Overall, we find that the greater degree prawns are generalist predators, the more a health-maximizing decision-maker will initially feed since it needs to compensate for the decreased effect of feed to get the same effect as when the prawn is more a specialist predator (Figure C.6; compare the initial level of the blue line in Panel A and Panel C).

3.4 Discussion

We show the importance of considering the impact of human behavior when designing public health policies. Our work contributes to the growing literature that predation of intermediate hosts can be an effective way to control schistosomiasis [42; 116; 123; 137; 138]; a debilitating parasitic disease that is transmitted to humans via contact with contaminated freshwater [35]. Previous studies have focused on identifying potential predators of the intermediate host (freshwater snails) [137], found that introducing non-native predators reduces disease prevalence in humans [123], showed that prawns can reduce snail-to-man transmission of schistosomiasis [116], and showed that prawn aquaculture reduces the disease's intermediate host population, and in turn aids in schistosomiasis control while being a source of sustainable development [42]. While these are important contributions towards a better understanding of how biological controls can help reduce schistosomiasis prevalence in the environment, our work contributes to this literature by showing how human behavior can potentially mitigate these public health benefits, and by offering potential solutions to this problem.

While previous work has taken into account the aquaculture farmers' profit-maximizing behavior (e.g. Hoover *et al.* [42] maximize profits over a 10-year planning horizon), they do not consider that profit-maximization may lead to changes in farming practices. In this paper, we allow for the possibility that aquaculture farmers can supplement the prawns' diet with feed. In turn, the objective of the decision-maker drives the intensity (i.e. the intensive margin) and extent (i.e. the extensive margin) to which feed is utilized, where the former represents the amount of feed provided to prawns in a given period and the latter represents the amount of time over which feed in provided (which in many cases is also the rotation length).

In our model, it is assumed that supplemental feed results in an increase growth of prawns and a reduction in the predation of snails. We find that if supplemental feeding results in the faster growth of prawns, the aquaculture farmer will be incentivized to feed the prawns as long as it is economical to do so (i.e., as long as the additional revenues

outweigh the feeding costs). If supplemental feeding results in a decreased predation of prawns on snails, our results suggest that the ancillary public health benefits of prawn aquaculture will be mitigated. The extent to which the public health benefits are reduced depends on several factors, including the prawns efficiency to convert supplemental feed into growth and the degree to which prawns are specialist or generalist predators. We find that a higher feed efficiency and a higher degree of generalist predation lead to decreased predation of snails, and in turn worsened human health outcomes.

Our results suggest that policy-makers and non-governmental organizations (NGOs) designing aquaculture programs in order to reduce schistosomiasis prevalence in the environment should account for the possibility that the profit-maximizing behavior of prawn aquaculture farmers may mitigate the anticipated public health benefits. Here, we propose two different policies that attenuate the effect of supplemental feeding: (i) a minimum, or standardized, rotation length and (ii) a limited feeding season. The standardized rotation length outperforms the limited feeding season when the health-maximizing rotation length (and thus minimum rotation length) is longer; this is the case if feed conversion efficiency is higher or when prawns are more generalist predators. The limited feeding season performs better when the health-maximizing rotation length (and thus the feeding season) is shorter; this occurs when feed conversion efficiency is low.

An important limitation of our paper is the assumption that human contact with freshwater only occurs within the boundaries of the enclosure and that the enclosure's snail and parasite populations are closed. Although relaxing these assumptions would dampen the negative impact of feed, we do believe that it would not change our main result: we need to account for potential unintended consequences that human behavior may have when designing such public health policies. Further modelling issues would also necessarily arise because assumptions, or further studies, would be required to understand how humans behave with the prawn enclosure and non-enclosed open water, and on how parasites and snails move in and out of the enclosure (see [139] for an example of how river currents can affect movement of cercariae).

While our paper makes important contributions to understanding what behavioral aspects need to be taken into account when designing and implementing large-scale aquaculture projects in order to fight schistosomiasis, it does not consider that policy-makers and NGOs implementing these policies are likely to also be responsible for other public health interventions meant to reduce schistosomiasis transmission, e.g., mass drug administrations of praziquantel [6]. This is important because there is mounting evidence that combining several types of treatments can make a considerable difference in the overall treatment cost and disease burden [5; 12; 19; 116]. Perhaps more importantly, integrated treatment guidelines for different types of treatment is necessary to avoid an over-utilization of public funds [100]. We believe future work should incorporate prawn aquaculture into integrated public health policies combining pharmaceutical and non-pharmaceutical interventions, rather than considering it as a separate and independent policy tool.

References

[1] Hotez PJ, Alvarado M, Basáñez MG, Bolliger I, Bourne R, Boussinesq M, et al. The global burden of disease study 2010: interpretation and implications for the neglected tropical diseases. PLoS neglected tropical diseases. 2014;8(7):e2865.

[2] Hotez PJ, Kamath A. Neglected tropical diseases in sub-Saharan Africa: review of their prevalence, distribution, and disease burden. PLoS Neglected Tropical Diseases. 2009;3(8):e412.

[3] World Health Organization. Accelerating work to overcome the global impact of neglected tropical diseases: a roadmap for implementation: executive summary. Geneva: World Health Organization; 2012.

[4] Garchitorena A, Sokolow S, Roche B, Ngonghala C, Jocque M, Lund A, et al. Disease ecology, health and the environment: a framework to account for ecological and socio-economic drivers in the control of neglected tropical diseases. Philosophical Transactions of the Royal Society B: Biological Sciences. 2017;372(1722):20160128.

[5] Sokolow SH, Wood CL, Jones IJ, Swartz SJ, Lopez M, Hsieh MH, et al. Global assessment of schistosomiasis control over the past century shows targeting the snail intermediate host works best. PLoS Neglected Tropical Diseases. 2016;10(7):e0004794.

[6] World Health Organization. Preventive chemotherapy in human helminthiasis. Coordinated use of anthelminthic drugs in control interventions: a manual for health professionals and programme managers. World Health Organization; 2006.

[7] Barnish G. Evaluation of Chemotherapy in the Control of Schistosoma Mansoni in Marquis Valley, Saint Lucia. The American Journal of Tropical Medicine and Hygiene. 1982;31(1):111–115.

[8] Inobaya MT, Olveda RM, Chau TN, Olveda DU, Ross AG. Prevention and control of schistosomiasis: a current perspective. Research and Reports in Tropical Medicine. 2014;2014(5):65.

[9] King CH, Olbrych SK, Soon M, Singer ME, Carter J, Colley DG. Utility of repeated praziquantel dosing in the treatment of schistosomiasis in high-risk communities in Africa: a systematic review. PLoS Neglected Tropical Diseases. 2011;5(9):e1321.

[10] Leslie J, Garba A, Oliva EB, Barkire A, Tinni AA, Djibo A, et al. Schistosomiais and soil-transmitted helminth control in Niger: cost effectiveness of school based and community distributed mass drug administration. PLoS Neglected Tropical Diseases. 2011;5(10):e1326.

[11] Lo NC, Bogoch II, Blackburn BG, Raso G, N'Goran EK, Coulibaly JT, et al. Comparison of community-wide, integrated mass drug administration strategies for schistosomiasis and soil-transmitted helminthiasis: a cost-effectiveness modelling study. The Lancet Global Health. 2015;3(10):e629–e638.

[12] Lo NC, Gurarie D, Yoon N, Coulibaly JT, Bendavid E, Andrews JR, et al. Impact and cost-effectiveness of snail control to achieve disease control targets for schistosomiasis. Proceedings of the National Academy of Sciences. 2018;115(4):E584–E591.

[13] Dyson L, Stolk WA, Farrell SH, Hollingsworth TD. Measuring and modelling the effects of systematic non-adherence to mass drug administration. Epidemics. 2017;18:56–66.

[14] Utzinger J, Zhou XN, Chen MG, Bergquist R. Conquering schistosomiasis in China: the long march. Acta Tropica. 2005;96(2-3):69–96.

[15] Guo JG, Chun-Li C, Guang-Han H, Han L, Dong L, Rong Z, et al. The role of 'passive chemotherapy'plus health education for schistosomiasis control in China during maintenance and consolidation phase. Acta Tropica. 2005;96(2-3):177–183.

[16] Ross A, Bartley P, Sleigh A, Olds G, Li Y, Williams G, et al. Schistosomiasis. New England Journal of Medicine. 2002;346:1212–1220.

[17] Hotez PJ, Fenwick A, Kjetland EF. Africa's 32 cents solution for HIV/AIDS. PLoS Neglected Tropical Diseases. 2009;3(5):e430.

[18] Andres L, Borja-Vega C, Fenwick C, de Jesus Filho J, Gomez-Suarez R. Overview and meta-analysis of global water, sanitation, and hygiene (WASH) impact evaluations. The World Bank; 2018.

[19] King CH, Sutherland LJ, Bertsch D. Systematic review and meta-analysis of the impact of chemical-based molluscicide for control of Schistosoma mansoni and S. haematobium transmission. PLoS Neglected Tropical Diseases. 2015;9(12):e0004290.

[20] Sokolow SH, Huttinger E, Jouanard N, Hsieh MH, Lafferty KD, Kuris AM, et al. Reduced transmission of human schistosomiasis after restoration of a native river prawn that preys on the snail intermediate host. Proceedings of the National Academy of Sciences. 2015;112(31):9650–9655.

[21] Ding C, Qiu Z, Zhu H. Multi-host transmission dynamics of schistosomiasis and its optimal control. Mathematical Biosciences & Engineering. 2015;12(5):983–1006.

[22] Ding C, Sun Y, Zhu Y. A schistosomiasis compartment model with incubation and its optimal control. Mathematical Methods in the Applied Sciences. 2017;40(14):5079–5094.

[23] Okosun KO, et al. Optimal control analysis of Malaria-Schistosomiasis co-infection dynamics. Mathematical Biosciences & Engineering. 2016;14(2):377.

[24] Kalinda C, Mushayabasa S, Chimbari MJ, Mukaratirwa S. Optimal control applied to a temperature dependent schistosomiasis model. Biosystems. 2019;175:47–56.

[25] Diekmann O, Heesterbeek JAP, Metz JA. On the definition and the computation of the basic reproduction ratio R0 in models for infectious diseases in heterogeneous populations. Journal of Mathematical Biology. 1990;28(4):365–382.

[26] Woolhouse M, Mutapi F, Ndhlovu P, Chandiwana S, Hagan P. Exposure, infection and immune responses to Schistosoma haematobium in young children. Parasitology. 2000;120(1):37–44.

[27] Goddard M, Jordan P. On the longevity of Schistosoma mansoni in man on St. Lucia, West Indies. Transactions of the Royal Society of Tropical Medicine and Hygiene. 1980;74(2):185–191.

[28] Anderson R, May R. Prevalence of schistosome infections within molluscan populations: observed patterns and theoretical predictions. Parasitology. 1979;79(1):63–94.

[29] Kariuki HC, Ivy JA, Muchiri EM, Sutherland LJ, King CH. Divergent effects of Schistosoma haematobium exposure on intermediate-host snail species Bulinus nasutus and Bulinus globosus from coastal Kenya. The American Journal of Tropical Medicine and Hygiene. 2017;96(4):850–855.

[30] Woolhouse M, Chandiwana S. Population biology of the freshwater snail Bulinus globosus in the Zimbabwe highveld. Journal of Applied Ecology. 1990:41–59.

[31] Central Intelligence Agency. The World Factbook. Washington, DC: Central Intelligence Agency, 2018. `https://www.cia.gov/library/publications/the-world-factbook/index.html` (accessed November 10, 2018); 2018.

[32] Guyatt H. The cost of delivering and sustaining a control programme for schistosomiasis and soil-transmitted helminthiasis. Acta Tropica. 2003;86(2-3):267–274.

[33] Barnish G, Jordan P, Bartholomew RK, Grist E. Routine focal molluscicide after chemotherapy to control Schistosoma mansoni in Cul de Sac valley, Saint Lucia. Transactions of the Royal Society of Tropical Medicine and Hygiene. 1982;76(5):602–609.

[34] Yang GJ, Li W, Sun LP, Wu F, Yang K, Huang YX, et al. Molluscicidal efficacies of different formulations of niclosamide: result of meta-analysis of Chinese literature. Parasites & Vectors. 2010;3(1):84.

[35] Colley DG, Bustinduy AL, Secor WE, King CH. Human schistosomiasis. The Lancet. 2014;383(9936):2253–2264.

[36] Miguel E, Kremer M. Worms: identifying impacts on education and health in the presence of treatment externalities. Econometrica. 2004;72(1):159–217.

[37] Ndamba J, Makaza N, Munjoma M, Gomo E, Kaondera K. The physical fitness and work performance of agricultural workers infected with Schistosoma mansoni in Zimbabwe. Annals of Tropical Medicine & Parasitology. 1993;87(6):553–561.

[38] Fenwick A, Figenschou B. The effect of Schistosoma mansoni infection on the productivity of cane cutters on a sugar estate in Tanzania. Bulletin of the World Health Organization. 1972;47(5):567.

[39] Kamel M, Moustafa Y, Foda N, Khashab S, Moemen M, Abo El Naga R. Impact of schistosomiasis on quality of life and productivity of workers. Eastern Mediterranean Health Journal. 2002;8(2-3):354–362.

[40] Weisbrod BA, Helminiak TW. Parasitic diseases and agricultural labor productivity. Economic Development and Cultural Change. 1977;25(3):505–522.

[41] Gurarie D, King CH, Yoon N, Li E. Refined stratified-worm-burden models that incorporate specific biological features of human and snail hosts provide better estimates of Schistosoma diagnosis, transmission, and control. Parasites & Vectors. 2016;9(1):428.

[42] Hoover CM, Sokolow SH, Kemp J, Sanchirico JN, Lund AJ, Jones IJ, et al. Modelled effects of prawn aquaculture on poverty alleviation and schistosomiasis control. Nature Sustainability. 2019;2(7):611.

[43] Spence M, Starrett D. Most rapid approach paths in accumulation problems. International Economic Review. 1975:388–403.

[44] Sanchirico JN, Springborn M. How to get there from here: ecological and economic dynamics of ecosystem service provision. Environmental and Resource Economics. 2011;48(2):243–267.

[45] Kling DM, Sanchirico JN, Wilen JE. Bioeconomics of managed relocation. Journal of the Association of Environmental and Resource Economists. 2016;3(4):1023–1059.

[46] Fuller KB, Sanchirico JN, Alston JM. The spatial-dynamic benefits from cooperative disease control in a perennial crop. Journal of Agricultural and Resource Economics. 2017;42(2):127–145.

[47] Rowthorn RE, Laxminarayan R, Gilligan CA. Optimal control of epidemics in metapopulations. Journal of the Royal Society Interface. 2009;6(41):1135–1144.

[48] Andrews P, Thyssen J, Lorke D. The biology and toxicology of molluscicides, Bayluscide. Pharmacology & Therapeutics. 1982;19(2):245–295.

[49] Oliveira-Filho EC, Paumgartten FJ. Toxicity of Euphorbia milii latex and niclosamide to snails and nontarget aquatic species. Ecotoxicology and Environmental Safety. 2000;46(3):342–350.

[50] Dai Jr, Wang W, Liang Ys, Li Hj, Guan Xh, Zhu Yc. A novel molluscicidal formulation of niclosamide. Parasitology Research. 2008;103(2):405–412.

[51] World Health Organization. Molluscicides. Second report of the Expert Committee on Bilharziasis. 1961.

[52] McCullough F, Gayral P, Duncan J, Christie J. Molluscicides in schistosomiasis control. Bulletin of the World Health Organization. 1980;58(5):681.

[53] McCullough FS, et al. The role of mollusciciding in schistosomiasis control. 1992.

[54] World Health Organization. Field use of molluscicides in schistosomiasis control programmes: an operational manual for programme managers. 2017.

[55] Thomas MB, Read AF. The threat (or not) of insecticide resistance for malaria control. Proceedings of the National Academy of Sciences. 2016;113(32):8900–8902.

[56] Sokolow SH, Nova N, Pepin KM, Peel AJ, Pulliam JR, Manlove K, et al. Ecological interventions to prevent and manage zoonotic pathogen spillover. Philosophical Transactions of the Royal Society B. 2019;374(1782):20180342.

[57] Mbah MLN, Poolman EM, Atkins KE, Orenstein EW, Meyers LA, Townsend JP, et al. Potential cost-effectiveness of schistosomiasis treatment for reducing HIV transmission in Africa–the case of Zimbabwean women. PLoS Neglected Tropical Diseases. 2013;7(8):e2346.

[58] Emanuel E, Persad G, Upshur R, Thome B, Parker M, Glickman A, et al. Fair Allocation of Scarce Medical Resources in the Time of COVID-19. The New England Journal of Medicine. 2020.

[59] Roope LS, Buckell J, Becker F, Candio P, Violato M, Sindelar JL, et al. How Should a Safe and Effective COVID-19 Vaccine be Allocated? Health Economists Need to be Ready to Take the Baton. PharmacoEconomics-Open. 2020:1–5.

[60] Buckner JH, Chowell G, Springborn MR. Dynamic Prioritization of COVID-19 Vaccines When Social Distancing is Limited for Essential Workers. medRxiv. 2020.

[61] Emanuel EJ, Persad G, Kern A, Buchanan A, Fabre C, Halliday D, et al. An ethical framework for global vaccine allocation. Science (New York, NY). 2020:eabe2803.

[62] Liu Y, Salwi S, Drolet B. Multivalue ethical framework for fair global allocation of a COVID-19 vaccine. Journal of Medical Ethics. 2020.

[63] Yamey G, Schäferhoff M, Hatchett R, Pate M, Zhao F, McDade KK. Ensuring global access to COVID-19 vaccines. The Lancet. 2020;395(10234):1405–1406.

[64] World Health Organization and others. WHO SAGE values framework for the allocation and prioritization of COVID-19 vaccination, 14 September 2020. World Health Organization; 2020.

[65] National Academies of Sciences, Engineering, and Medicine. Discussion Draft of the Preliminary Framework for Equitable Allocation of COVID-19 Vaccine. 2020. https://doi.org/10.17226/25914.

[66] Zaric GS, Brandeau ML. Optimal investment in a portfolio of HIV prevention programs. Medical Decision Making. 2001;21(5):391–408.

[67] Verity R, Okell LC, Dorigatti I, Winskill P, Whittaker C, Imai N, et al. Estimates of the severity of coronavirus disease 2019: a model-based analysis. The Lancet infectious diseases. 2020.

[68] Baker MG, Peckham TK, Seixas NS. Estimating the burden of United States workers exposed to infection or disease: a key factor in containing risk of COVID-19 infection. PLoS One. 2020;15(4):e0232452.

[69] Thomas LJ, Huang P, Yin F, Luo XI, Almquist ZW, Hipp JR, et al. Spatial heterogeneity can lead to substantial local variations in COVID-19 timing and severity. Proceedings of the National Academy of Sciences. 2020;117(39):24180–24187.

[70] Polyakova M, Kocks G, Udalova V, Finkelstein A. Initial economic damage from the COVID-19 pandemic in the United States is more widespread across ages and geographies than initial mortality impacts. Proceedings of the National Academy of Sciences. 2020. Available from: https://www.pnas.org/content/early/2020/10/19/2014279117.

[71] Brandeau ML, Zaric GS, Richter A. Resource allocation for control of infectious diseases in multiple independent populations: beyond cost-effectiveness analysis. Journal of health economics. 2003;22(4):575–598.

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

[73] Gersovitz M, Hammer JS. The economical control of infectious diseases. The Economic Journal. 2004;114(492):1–27.

[74] Dangerfield CE, Vyska M, Gilligan CA. Resource Allocation for Epidemic Control Across Multiple Sub-populations. Bulletin of mathematical biology. 2019;81(6):1731–1759.

[75] Acemoglu D, Chernozhukov V, Werning I, Whinston MD. Optimal targeted lockdowns in a multi-group SIR model. NBER Working Paper. 2020;27102.

[76] Alvarez FE, Argente D, Lippi F. A simple planning problem for COVID-19 lockdown. NBER Working Paper. 2020;26981.

[77] Levine-Tiefenbrun M, Yelin I, Katz R, Herzel E, Golan Z, Schreiber L, et al. Decreased SARS-CoV-2 viral load following vaccination. medRxiv. 2021.

[78] Begon M, Bennett M, Bowers RG, French NP, Hazel S, Turner J. A clarification of transmission terms in host-microparasite models: numbers, densities and areas. Epidemiology & Infection. 2002;129(1):147–153.

[79] Edridge AW, Kaczorowska J, Hoste AC, Bakker M, Klein M, Loens K, et al. Seasonal coronavirus protective immunity is short-lasting. Nature medicine. 2020;26(11):1691–1693.

[80] Peng L, Yang W, Zhang D, Zhuge C, Hong L. Epidemic analysis of COVID-19 in China by dynamical modeling. arXiv preprint arXiv:200206563. 2020.

[81] Yang Z, Zeng Z, Wang K, Wong SS, Liang W, Zanin M, et al. Modified SEIR and AI prediction of the epidemics trend of COVID-19 in China under public health interventions. Journal of Thoracic Disease. 2020;12(3):165.

[82] Hou C, Chen J, Zhou Y, Hua L, Yuan J, He S, et al. The effectiveness of quarantine of Wuhan city against the Corona Virus Disease 2019 (COVID-19): A well-mixed SEIR model analysis. Journal of medical virology. 2020.

[83] Pandey G, Chaudhary P, Gupta R, Pal S. SEIR and Regression Model based COVID-19 outbreak predictions in India. arXiv preprint arXiv:200400958. 2020.

[84] Prem K, Liu Y, Russell TW, Kucharski AJ, Eggo RM, Davies N, et al. The effect of control strategies to reduce social mixing on outcomes of the COVID-19 epidemic in Wuhan, China: a modelling study. The Lancet Public Health. 2020.

[85] Radulescu A, Cavanagh K. Management strategies in a SEIR model of COVID 19 community spread. arXiv preprint arXiv:200311150. 2020.

[86] Roda WC, Varughese MB, Han D, Li MY. Why is it difficult to accurately predict the COVID-19 epidemic? Infectious Disease Modelling. 2020.

[87] Bjørnstad ON, Shea K, Krzywinski M, Altman N. Modeling infectious epidemics. Nature methods. 2020;17:455–456.

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

[89] Stutt RO, Retkute R, Bradley M, Gilligan CA, Colvin J. A modelling framework to assess the likely effectiveness of facemasks in combination with 'lockdown'in managing the COVID-19 pandemic. Proceedings of the Royal Society A. 2020;476(2238):20200376.

[90] Bertozzi AL, Franco E, Mohler G, Short MB, Sledge D. The challenges of modeling and forecasting the spread of COVID-19. Proceedings of the National Academy of Sciences. 2020;117(29):16732–16738.

[91] Burton J, Billings L, Cummings DA, Schwartz IB. Disease persistence in epidemiological models: the interplay between vaccination and migration. Mathematical biosciences. 2012;239(1):91–96.

[92] Chen M, Li M, Hao Y, Liu Z, Hu L, Wang L. The introduction of population migration to SEIAR for COVID-19 epidemic modeling with an efficient intervention strategy. Information Fusion. 2020;64:252–258.

[93] Goldman SM, Lightwood J. Cost optimization in the SIS model of infectious disease with treatment. The BE Journal of Economic Analysis & Policy. 2002;2(1).

[94] Barrett S, Hoel M. Optimal disease eradication. Environment and Development Economics. 2007:627–652.

[95] Sivaraman NK, Gaur M, Baijal S, Rupesh CV, Muthiah SB, Sheth A. Exo-SIR: An Epidemiological Model to Analyze the Impact of Exogenous Infection of COVID-19 in India. arXiv preprint arXiv:200806335. 2020.

[96] World Tourism Organization. International Tourism Down 70% As Travel Restrictions Impact All Regions. Madrid, Spain: United Nations World Tourism Organization, 2020. https://www.unwto.org/news/international-tourism-down-70-as-travel-restrictions-impact-all-regions (accessed October 27, 2020); 2020.

[97] Environmental Protection Agency. What value of statistical life does EPA use? Washington, DC: Environmental Protection Agency, 2020. https://www.epa.gov/environmental-economics/mortality-risk-valuation#whatvalue (accessed October 27, 2020); 2020.

[98] Nurchis MC, Pascucci D, Sapienza M, Villani L, D'Ambrosio F, Castrini F, et al. Impact of the Burden of COVID-19 in Italy: Results of Disability-Adjusted Life Years (DALYs) and Productivity Loss. International Journal of Environmental Research and Public Health. 2020;17(12):4233.

[99] Ryan D, Toews C, Sanchirico JN, Armsworth PR. Implications of policy adjustment costs for fisheries management. Natural Resource Modeling. 2017;30(1):74–90.

[100] Castonguay FM, Sokolow SH, De Leo GA, Sanchirico JN. Cost-effectiveness of combining drug and environmental treatments for environmentally transmitted diseases. Proceedings of the Royal Society B. 2020;287(1933):20200966.

[101] Ohmit SE, Thompson MG, Petrie JG, Thaker SN, Jackson ML, Belongia EA, et al. Influenza vaccine effectiveness in the 2011–2012 season: protection against each circulating virus and the effect of prior vaccination on estimates. Clinical infectious diseases. 2014;58(3):319–327.

[102] Duijzer LE, van Jaarsveld WL, Wallinga J, Dekker R. Dose-optimal vaccine allocation over multiple populations. Production and Operations Management. 2018;27(1):143–159.

[103] Phadke I, McKee A, Conway JM, Shea K. Analysing how changes in the health status of healthcare workers affects epidemic outcomes. Epidemiology and Infection. 2021;149:e42.

[104] Verelst F, Kuylen E, Beutels P. Indications for healthcare surge capacity in European countries facing an exponential increase in coronavirus disease (COVID-19) cases, March 2020. Eurosurveillance. 2020;25(13):2000323.

[105] Baker SR, Farrokhnia RA, Meyer S, Pagel M, Yannelis C. How does household spending respond to an epidemic? consumption during the 2020 COVID-19 pandemic. National Bureau of Economic Research; 2020.

[106] Andersen AL, Hansen ET, Johannesen N, Sheridan A. Consumer responses to the COVID-19 crisis: Evidence from bank account transaction data. Available at SSRN 3609814. 2020.

[107] Coibion O, Gorodnichenko Y, Weber M. The cost of the COVID-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. National Bureau of Economic Research; 2020.

[108] Vestergaard LS, Nielsen J, Richter L, Schmid D, Bustos N, Braeye T, et al. Excess all-cause mortality during the COVID-19 pandemic in Europe–preliminary pooled estimates from the EuroMOMO network, March to April 2020. Eurosurveillance. 2020;25(26):2001214.

[109] Pfefferbaum B, North CS. Mental health and the Covid-19 pandemic. New England Journal of Medicine. 2020.

[110] Ssentongo P, Ssentongo AE, Heilbrunn ES, Ba DM, Chinchilli VM. Association of cardiovascular disease and 10 other pre-existing comorbidities with COVID-19 mortality: A systematic review and meta-analysis. PloS one. 2020;15(8):e0238215.

[111] Hu B, Guo H, Zhou P, Shi ZL. Characteristics of SARS-CoV-2 and COVID-19. Nature Reviews Microbiology. 2020:1–14.

[112] Fenner F, Henderson DA, Arita I, Jezek Z, Ladnyi ID, et al. Smallpox and its eradication. vol. 6. Geneva: World Health Organization; 1988.

[113] Rossman H, Shilo S, Meir T, Gorfine M, Shalit U, Segal E. COVID-19 dynamics after a national immunization program in Israel. Nature Medicine. 2021:1–7.

[114] Haas EJ, Angulo FJ, McLaughlin JM, Anis E, Singer SR, Khan F, et al. Impact and effectiveness of mRNA BNT162b2 vaccine against SARS-CoV-2 infections and COVID-19 cases, hospitalisations, and deaths following a nationwide vaccination campaign in Israel: an observational study using national surveillance data. The Lancet. 2021.

[115] McManus D, Dunne D, Sacko M, Utzinger J, Vennervald B, Zhou X. Schistosomiasis. Nat Rev Dis Primers. 2018;4(1):13.

[116] Sokolow SH, Huttinger E, Jouanard N, Hsieh MH, Lafferty KD, Kuris AM, et al. Reduced transmission of human schistosomiasis after restoration of a native river prawn that preys on the snail intermediate host. Proceedings of the National Academy of Sciences. 2015;112(31):9650–9655.

[117] Hammond A, Galizi R, Kyrou K, Simoni A, Siniscalchi C, Katsanos D, et al. A CRISPR-Cas9 gene drive system targeting female reproduction in the malaria mosquito vector Anopheles gambiae. Nature biotechnology. 2016;34(1):78–83.

[118] Sokolow SH, Wood CL, Jones IJ, Lafferty KD, Kuris AM, Hsieh MH, et al. To reduce the global burden of human schistosomiasis, use 'old fashioned'snail control. Trends in parasitology. 2018;34(1):23–40.

[119] Tanaka H, Tsuji M. From discovery to eradication of schistosomiasis in Japan: 1847–1996. International journal for parasitology. 1997;27(12):1465–1480.

[120] Laamrani H, Khallaayoune K, Boelee E, Laghroubi M, Madsen H, Gryseels B. Evaluation of environmental methods to control snails in an irrigation system in Central Morocco. Tropical Medicine & International Health. 2000;5(8):545–552.

[121] Incani RN. The Venezuelan experience in the control of schistosomiasis mansoni. Mem Inst Oswaldo Cruz. 1987;82(Suppl IV):89–93.

[122] Al-Madani A. Schistosomiasis control in Saudi Arabia with special reference to the period 1983–1988. Public Health. 1990;104(4):261–266.

[123] Mkoji GM, Hofkin BV, Kuris AM, Stewart-Oaten A, Mungai BN, Kihara JH, et al. Impact of the crayfish Procambarus clarkii on Schistosoma haematobium transmission in Kenya. The American journal of tropical medicine and hygiene. 1999;61(5):751–759.

[124] Howarth FG. Environmental impacts of classical biological control. Annual review of entomology. 1991;36(1):485–509.

[125] Bale J, Van Lenteren J, Bigler F. Biological control and sustainable food production. Philosophical Transactions of the Royal Society B: Biological Sciences. 2008;363(1492):761–776.

[126] Bowes MD, Krutilla JV. Multiple use management of public forestlands. In: Handbook of natural resource and energy economics. vol. 2. Elsevier; 1985. p. 531–569.

[127] Lima JdF, Garcia JdS, Silva TCd. Natural diet and feeding habits of a freshwater prawn (Macrobrachium carcinus: Crustacea, Decapoda) in the estuary of the Amazon River. Acta Amazonica. 2014;44(2):235–244.

[128] Ranjeet K, Kurup B. Heterogeneous individual growth of Macrobrachium rosenbergii male morphotypes. Naga, The ICLARM Quarterly. 2002;25(2):13–18.

[129] Sokolow SH, Lafferty KD, Kuris AM. Regulation of laboratory populations of snails (Biomphalaria and Bulinus spp.) by river prawns, Macrobrachium

spp.(Decapoda, Palaemonidae): implications for control of schistosomiasis. Acta tropica. 2014;132:64–74.

[130] Lorenzen K. The relationship between body weight and natural mortality in juvenile and adult fish: a comparison of natural ecosystems and aquaculture. Journal of fish biology. 1996;49(4):627–642.

[131] Ventura T, Sagi A. The insulin-like androgenic gland hormone in crustaceans: From a single gene silencing to a wide array of sexual manipulation-based biotechnologies. Biotechnology Advances. 2012;30(6):1543–1550.

[132] Aflalo ED, Hoang TTT, Nguyen V, Lam Q, Nguyen D, Trinh Q, et al. A novel two-step procedure for mass production of all-male populations of the giant freshwater prawn Macrobrachium rosenbergii. Aquaculture. 2006;256(1-4):468–478.

[133] Levy T, Rosen O, Eilam B, Azulay D, Aflalo ED, Manor R, et al. A single injection of hypertrophied androgenic gland cells produces all-female aquaculture. Marine Biotechnology. 2016;18(5):554–563.

[134] Holling CS. The components of predation as revealed by a study of small-mammal predation of the European pine sawfly. The Canadian Entomologist. 1959;91(5):293–320.

[135] Halstead NT, Hoover CM, Arakala A, Civitello DJ, De Leo GA, Gambhir M, et al. Agrochemicals increase risk of human schistosomiasis by supporting higher densities of intermediate hosts. Nature communications. 2018;9(1):837.

[136] Castonguay FM, Blackwood JC, Howerton E, Shea K, Sims C, Sanchirico JN. Spatial Allocation of Scarce COVID-19 Vaccines. medRxiv. 2021. Available from: `https://www.medrxiv.org/content/early/2021/03/14/2020.12.18.20248439`.

[137] Younes A, El-Sherief H, Gawish F, Mahmoud M. Biological control of snail hosts transmitting schistosomiasis by the water bug, Sphaerodema urinator. Parasitology research. 2017;116(4):1257–1264.

[138] Alkalay AS, Rosen O, Sokolow SH, Faye YP, Faye DS, Aflalo ED, et al. The prawn Macrobrachium vollenhovenii in the Senegal River basin: towards sustainable restocking of all-male populations for biological control of schistosomiasis. PLoS Negl Trop Dis. 2014;8(8):e3060.

[139] Ciddio M, Mari L, Sokolow SH, De Leo GA, Casagrandi R, Gatto M. The spatial spread of schistosomiasis: A multidimensional network model applied to Saint-Louis region, Senegal. Advances in water resources. 2017;108:406–415.

Appendices

Contents

A Chapter 1

A.1 Model Description

A.1.1 Epidemiological Model

We assume two closed human sub-populations $i = C, A$, representing the children (C) and adult (A) populations, respectively. Let I_i denote the proportion of infected humans in population i. We incorporate into a single parameter the rate of reproduction of propagules released in the environment by population i and the fraction of propagules that reach the environment; this parameter represents the shedding rate of infectious humans and it is denoted by λ_i.

We separate the population of intermediate hosts (i.e., the freshwater snails) in two epidemiological classes: susceptible (X) and infected (W). We denote the total population of intermediate hosts as N which means that by definition $0 \leq X + W = N \leq K$ where K is the carrying capacity of the ecosystem. Following the convention in the literature on disease dynamics [1; 2; 3; 4; 5; 6; 7], we model the growth of infected intermediate hosts as

$$\dot{W} = \Big(\alpha_C \lambda_C I_C + \alpha_A \lambda_A I_A\Big) X - \Big(\delta + q_W u_W\Big) W \tag{A.1}$$

where α_i represents the proportion of population i in total population, δ is the natural death rate of infected intermediate hosts, u_W is the level of environmental treatment (EnvTre), and q_W, with $0 < q_W < 1$, represents the fact that there is an upper bound on the effectiveness of EnvTre [8]. In this formulation, entry of intermediate hosts in the infectious compartment is represented by the first term of Equation (A.1), and exit of intermediate hosts from the infectious compartment (i.e., death of infected intermediate hosts) is represented by the second term of Equation (A.1). New infections of intermediate hosts thus rely on the relative (α_i) shedding (λ_i) of infectious humans (I_i), while the loss of infectious intermediate hosts relies on the natural death rate (δ) and the fact that we apply an imperfect (q_W) and non-selective treatment (u_W) that kills both susceptible and infected intermediate hosts.

The intermediate hosts are assumed to grow according to a logistic-type reproduction function [2; 3; 5]. Normalizing the carrying capacity K to one such that $0 \leq X + W = N \leq K = 1$, the growth of the susceptible class of intermediate hosts is

$$\dot{X} = fX\Big(1 - N\Big) - q_W u_W X - \Big(\alpha_C \lambda_C I_C + \alpha_A \lambda_A I_A\Big)X \tag{A.2}$$

where f is the maximal reproduction rate. Note that density dependent mortality (or fecundity) of susceptible snails is implicitly accounted for in the nonlinear term of the logistic growth function (i.e., accounted for in the first term of Equation A.2); only susceptible snails can reproduce since the parasite castrates the snails [3; 5]. Given that we model the control of intermediate hosts as non-selective, controlling the population of infected intermediate hosts also affects susceptible intermediate hosts (second term of Equation A.2). Finally, the third term of Equation (A.2) represents the susceptible snails becoming infected.

We focus on a case where there is contact with the disease exclusively through the environment; it infects susceptible humans at rate β_i (for $i = C, A$) through the W infected intermediate hosts living in the environmental reservoir. We assume that only children are treated via drugs, which allows us to focus on a more realistic case where only part of the human-to-environment transmission is reduced. The proportion of infected human children and adults respectively grow according to

$$\dot{I}_C = \beta_C W\Big(1 - I_C\Big) - \Big(\gamma_C + q_C u_C\Big)I_C \tag{A.3a}$$

$$\dot{I}_A = \beta_A W\Big(1 - I_A\Big) - \gamma_A I_A \tag{A.3b}$$

where γ_i (for $i = C, A$) is the natural rate of recovery of population i, u_C is the level of MDA treatment for children, and q_C, with $0 < q_C < 1$, represents the fact that there is an upper bound on the effectiveness of MDA controls [9; 10]. Note that because I_i represents a proportion of the population $i = C, A$, then $(1 - I_i)$ represents the proportion

of susceptible individuals since there is no immunity against schistosomiasis. See figure A.1 for a flow diagram of the disease transmission.

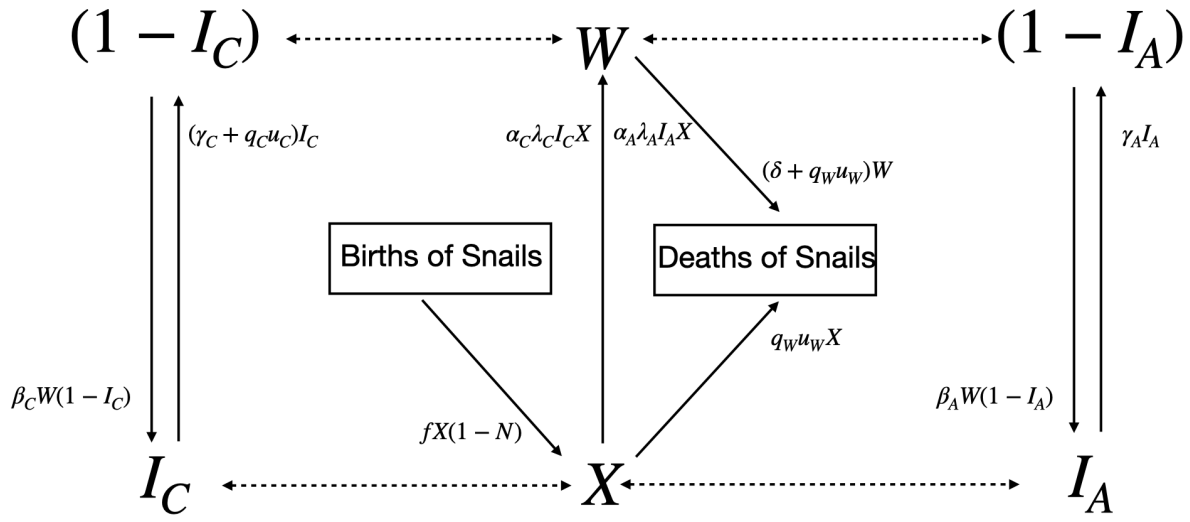

Figure A.1: This figure shows a flow diagram of the disease transmission. Full lines represent flows between state variables; dashed lines represents indirect infection between susceptible humans and infected snails, and infected humans and susceptible snails.

A.1.2 Economic Model

Damages due to infection occur on human populations $i = C, A$. We assume that the damage function is linear and additively separable such that for each population i the damage function is given by:

$$c_i(I_i) = c_i I_i \tag{A.4}$$

where c_i is the cost parameter associated with population i.

The cost of the policy is denoted by some function $c_P(u_C, u_W; \mathcal{F})$, where $\mathcal{F}$ represents a fixed transportation and management cost that is incurred each period during the planning period. We model it as being additively separable functions, such that

$$c_P(u_C, u_W; \mathcal{F}) = c_{PC}(u_C) + c_{PW}(u_W) + \mathcal{F}. \tag{A.5}$$

We assume that the cost of the policy for MDA treatment, $c_{PC}(u_C)$, and the cost of the EnvTre policy, $c_{PW}(u_W)$, are increasing in the level of control. The functional form chosen for treatment of children is given by,

$$c_{PC}(u_C) = c_{PC}u_C$$

where c_{PC} is the cost parameter associated with the MDA cost function of the treatment of children; this parameter can be thought of as the cost of the drug multiplied by the number of children in the population. We assume that costs associated with control of intermediate hosts is given by the following function:

$$c_{PW}(u_W) = c_{PW}u_W$$

where c_{PW} denotes the cost parameter associated with the treatment of intermediate hosts.

We model transportation and management costs, $\mathcal{F}$, as a fixed cost in each period during the planning period regardless of whether treatment is being undertaken. There are a number of ways to consider the fixed cost. We could assume a separate and additive fixed cost for each treatment. This implies, however, that the planning agency is sending separate simultaneous shipments and teams to the village, which seems unrealistic. On the other hand, there are likely potential economies of scale across the different treatment options. For example, an agency could combine transportation of treatments to a remote village. To account for this realistic possibility, we model a single fixed cost regardless of whether children or environmental treatments are applied. We parameterize the fixed cost from the literature (see, for example, [4; 5; 11; 12; 13; 14; 15]).

A.2 Parameterization

A.2.1 Epidemiological Model

In epidemiology, the basic reproduction ratio R_0 is defined as being the expected number of secondary infection, at a disease-free equilibrium, caused by a typical infected individual over its entire infectious period [16]. In a basic model, the R_0 is given by the contact rate multiplied by the mean infectious period. In more complex heterogeneous models however, one needs to use the next-generation matrix. The dominant eigenvalue of the next-generation matrix is the basic reproduction ratio R_0 [16], which turns out to be a function of relative shedding rates of adult and children, natural recovery rates, and contact rates with the environmental reservoir. The next-generation matrix is composed of two matrices, denoted F and V, and it is equal to $-FV^{-1}$. 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 [7]. Hence in our model,

$$F = \begin{pmatrix} 0 & 0 & \beta_C \\ 0 & 0 & \beta_A \\ 0 & 0 & 0 \end{pmatrix} \text{ and } V = \begin{pmatrix} -\gamma_C & 0 & 0 \\ 0 & -\gamma_A & 0 \\ \alpha_C\lambda_C & \alpha_A\lambda_A & -\delta \end{pmatrix}$$

where the three rows of F and V refer to the I_C, I_A, and W equations, respectively. Note that both F and V are derived under the assumption of introducing a single infected snail in an otherwise susceptible population.

Following the methodology of Diekmann *et al.* [16], the basic reproduction ratio is given by,

$$R_0 = \frac{\alpha_C\beta_C\lambda_C}{\delta\gamma_C} + \frac{\alpha_A\beta_A\lambda_A}{\delta\gamma_A}.$$

There is little evidence on how and to what extent the biological parameters in the above R_0 may differ between children and adults.[1] This evidence may be anecdotal or based on beliefs [4], and also may be highly community and occupation dependent [18]. We thus assume in our numerical simulation that all biological parameters are identical across the child and adult populations.

There is little data available on the contact rate of humans with the pathogen, β_i, and we found no data available on the shedding rate λ_i. In our simulation the contact rate is set such that there is 1 infection per 200 water contacts,[2] and the shedding rate is set such that there is approximately 1 intermediate host infection per 555 sheds. The chosen values of β_i and λ_i were calibrated to match a basic reproduction ratio, R_0, of 3.5.[3]

The natural recovery rate in humans, γ_i, may be thought of as the life expectancy of the disease in hosts (3.3 years) [20]. Since we do not take into consideration the intensity of infection human hosts, we assume that humans, once infected, carry 70 worms [3]. The maximum reproduction rate of intermediate hosts, f, is set according to estimates of the literature [21], assuming a carrying capacity of 10,000 individuals in a water access point of 200 square meters [3]. The natural mortality rate of intermediate hosts, δ, is based on their life expectancy (2 months) [21; 22; 23], again assuming a carrying capacity of 10,000 individuals in water access point of 200 square meters [3].

A.2.2 Economic Model

The costs include damages (reduction in human health), treatment (MDA and EnvTre), and a fixed transportation (from a central planning agency to an endemic remote village) and management cost. All costs are discounted using a discount rate of $r = 0.04$ or 4%.

Damage related to infectious humans were calibrated such that, without intervention, infection rates of 38% in a community of 5,000 people would yield losses of 550 disability-adjusted life years (DALYs) [5]. The value of a life year was set to be approximately the

[1] See Colley *et al.* [17] for differences in intensities of infection with age for schistosomiasis; these differences may be attributable to antiparasite immunity rather than reduced contact.

[2] This is consistent with findings of the literature [3; 19].

[3] According to Sokolow *et al.* [3], the expected R_0 for schistosomiasis ranges from 1 to 7.

median value of the GDP per capita of an African country (approximately $3,000 USD). Although we assume that costs are identical for children and adults, cost parameters—representing damages on the whole subpopulation and not just one individual—must differ if proportions of children and adults are not the same.

For the MDA cost function, we calibrated the cost function using the cost of the drug used to fight schistosomiasis [4; 5; 11; 13; 14]. The MDA cost parameters were calibrated for a 5,000 people community. Effectiveness of MDA treatment was assumed to be 80% [9; 10].

For the EnvTre cost function, we calibrated the linear term of the cost function using estimates for variable costs of snail control (e.g., chemical, personnel compensation) [5; 12; 15]. Chemical applications of molluscicide niclosamide does not systematically kill all snails; the meta-analysis of Yang *et al.* [8] finds that the death rate of snails 15 days after the spraying is approximately 88%.

Finally, the fixed cost of transportation and management is based on estimates from the literature [4; 5; 11; 12; 13; 14; 15].

A.2.3 Parameter Levels

Table A.1 summarizes the parameter values we used in our numerical simulation.

Parameters	Level	Definition
β_i	5.00×10^{-3}	Contact rate for $i = C, A$ (year^{-1}).[4]
λ_i	1.80×10^{-3}	Shedding rate for $i = C, A$ (year^{-1}).[5]
γ_i	4.30×10^{-3}	Natural recovery rate for $i = C, A$ (year^{-1}).[6]
δ	6.00×10^{-4}	Natural death rate of the disease in the environment (year^{-1}).[7]
f	1.168	Maximum reproduction rate of intermediate hosts (year^{-1}).[8]
α_1	0.4	Proportion of children in total population.[9]
α_2	0.6	Proportion of adults in total population.[9]
r	0.04	Annual discount rate.[10]
c_C	208,268	Damages related to infectious children (US Dollars).[11]
c_A	312,402	Damages related to infectious adults (US Dollars).[11]
c_{PC}	210	Cost of drug for treating children population (MDA) (US Dollars).[12]
c_{PW}	370	Cost of chemical treatment (EnvTre) (US Dollars).[13]
$\mathcal{F}$	1,500	Fixed transportation and management costs (US Dollars).[14]
q_C	0.8	Effectiveness of MDA control.[15]
q_W	0.88	Effectiveness of EnvTre control.[16]

Table A.1: Parameter levels used in the base case of the numerical simulation.

A.3 Optimization

A.4 Boundary Conditions and Constraints

The initial and terminal conditions of the four state variables are such that:

$$W(0) \text{ is given, and } W(T) \text{ is free;} \tag{A.6a}$$

$$X(0) \text{ is given, and } X(T) \text{ is free;} \tag{A.6b}$$

$$I_i(0) \text{ are given, and } I_i(T) \text{ are free for } i = C, A.^{17} \tag{A.6c}$$

We present results from a numerical simulation where initially all state variables are at their pre-treatment equilibria.[18] The chosen parameter values imply that before any treatment, the level of infection for both the children and adult populations has converged

[4]Calibrated to match a R_0 of 3.5. Set such that there is 1 infection per 200 water contacts (consistent with the literature) [3; 19]. See electronic supplementary material section A.5.1 for a sensitivity analysis of the results when varying the contact rate relative to the shedding rate while keeping R_0 constant.

[5]Calibrated to match a R_0 of 3.5. This gives approximately 1 environmental infection per 555 shedding. See electronic supplementary material section A.5.1 for a sensitivity analysis of the results when varying the contact rate relative to the shedding rate while keeping R_0 constant.

[6]Based on the life expectancy of the disease in hosts (3.3 years) [20], assuming humans carry 70 worms once infected [3].

[7]Based on the life expectancy of infected intermediate hosts (2 months) [21; 22; 23], assuming a carrying capacity of 10,000 individuals in a water access point of 200 square meters [3].

[8]Based on an instantaneous intrinsic fertility rate of snails of 0.16 per day per square meter, and assuming a carrying capacity of 10,000 individuals in a water access point of 200 square meters [3]. See electronic supplementary material section A.5.7 for a sensitivity analysis of the maximum reproduction rate of intermediate hosts.

[9]Based on data from Senegal from "The World Factbook" [24]. For a sensitivity analysis of the results while varying the proportion of children in the total population, see electronic supplementary material section A.5.2.

[10]See electronic supplementary material section A.5.4 for a sensitivity analysis of the results while varying discount rate.

[11]Based on an estimate of the number of DALYs lost when no intervention takes place over a ten year period [5]. See electronic supplementary material section A.5.5 for a sensitivity analysis of the value of a DALY.

[12]Based on the literature [4; 5; 11; 13; 14].

[13]Based on the literature [5; 12; 15].

[14]Approximate value based on the literature [4; 5; 11; 12; 13; 14; 15].

[15]Based on the literature [9; 10]. See electronic supplementary material section A.5.6 for a sensitivity analysis of the effectiveness of MDA control.

[16]Based on the literature [8]. See electronic supplementary material section A.5.6 for a sensitivity analysis of the effectiveness of environmental control.

[17]For WHO's guidelines on treatment and long-term control and eradication objectives for neglected tropical diseases see [11] and [25]; for guidelines on vector management see [26].

[18]See electronic supplementary material section A.5.3 for sensitivity analyses of the results when varying initial level of infection.

to approximately 38%.[19] The steady-state level of the snail population will converge to the carrying capacity (i.e. $K = 1$), while the number of infected snails will converge to 54%. These are the initial values of the state variables in our numerical simulation.

All terminal values on state variables are free. This means that the algorithm approximating the optimal control problem will optimally choose the infection prevalence in children, as well as the optimal number of infected snails at the end of the time horizon. This requirement of choosing optimal levels is however conditional on the fact that these optimal levels can be reached in the given time horizon.

State variables W, X, and I_i for $i = C, A$, are also subject to constraints that bound them from above and below. Formally, the constraints are given by:

$$0 \leq W \leq N \leq 1; \tag{A.7a}$$

$$0 \leq X \leq N \leq 1; \tag{A.7b}$$

$$0 \leq I_i \leq 1 \quad \text{for } i = C, A. \tag{A.7c}$$

since the carrying capacity of the ecosystem where the intermediate host resides has been normalized to one, and since I_i denotes the proportion of infected individuals.

Control variables u_W and u_C are modeled in a additive way. The former acts as an increase in the death rate of the intermediate host, while the latter acts as an increase in the rate of recovery of children. Control variables are subject to:

$$0 \leq u_W \leq 1; \tag{A.8a}$$

$$0 \leq u_C \leq 0.9. \tag{A.8b}$$

The upper-bound of 0.9 on the MDA control variable u_C represent the fact that there is systematic noncompliance to drug treatment [27; 28; 29].

[19] This is consistent with the findings of Lo *et al.* [5].

A.4.1 Time Horizon and Objective Function

Since we do not impose conditions on state variables at the end of the planing period, i.e., since the terminal conditions are free, the minimization of the objective guarantees that the terminal conditions are optimized conditional on being reachable in a finite number of periods, T. We assume that

$$T \text{ is given,} \tag{A.9}$$

rather than being optimaly chosen. Following previous literature investigating the cost-effectiveness of schistosomiasis [2; 5; 7; 30], the finite amount of time chosen for the simulation is $T = 10$ years. We consider longer time horizon in section A.5.8 of the electronic supplementary material.

We model the problem as a cost minimization. Given some non-negative discount rate r, the objective of the planner is to minimize the present discounted costs of the disease—which includes both damages and treatment costs—which is given by:

$$\min_{u_C, u_W} \int_0^T e^{-rt} \Big\{ c_C(I_C) + c_A(I_A) + c_P(u_C, u_W; \mathcal{F}) \Big\} \mathrm{d}t \tag{A.10}$$

$$\text{s.t. } (A.1), (A.2), (A.3), (A.6), (A.7), (A.8) \text{ and } (A.9).$$

where $c_i(I_i)$ for $i = C, A$ and $c_P(u_C, u_W; \mathcal{F})$ are respectively given by (A.4) and (A.5). Explicitly, the problem is

$$\min_{u_C, u_W} \int_0^T e^{-rt}\Big\{c_C I_C + c_A I_A + c_{PC} u_C + c_{PW} u_W + \mathcal{F}\Big\}\mathrm{d}t \tag{A.10}$$

$$\text{s.t. } \dot{W} = \Big(\alpha_C \lambda_C I_C + \alpha_A \lambda_A I_A\Big)X - \Big(\delta + q_W u_W\Big)W \tag{A.1}$$

$$\dot{X} = fX\Big(1 - N\Big) - q_W u_W X - \Big(\alpha_C \lambda_C I_C + \alpha_A \lambda_A I_A\Big)X \tag{A.2}$$

$$\dot{I}_C = \beta_C W\Big(1 - I_C\Big) - \Big(\gamma_C + q_C u_C\Big)I_C \tag{A.3a}$$

$$\dot{I}_A = \beta_A W\Big(1 - I_A\Big) - \gamma_A I_A \tag{A.3b}$$

$$W(0) \text{ is given and } W(T) \text{ is free;} \tag{A.6a}$$

$$X(0) \text{ is given and } X(T) \text{ is free;} \tag{A.6b}$$

$$I_i(0) \text{ are given and } I_i(T) \text{ are free for } i = C, A; \tag{A.6c}$$

$$0 \leq W \leq N \leq 1; \tag{A.7a}$$

$$0 \leq X \leq N \leq 1; \tag{A.7b}$$

$$0 \leq I_i \leq 1 \text{ for } i = C, A; \tag{A.7c}$$

$$0 \leq u_W \leq 1; \tag{A.8a*}$$

$$0 \leq u_C \leq 0.9; \tag{A.8b*}$$

$$T \text{ is given.} \tag{A.9}$$

where constraints (A.8) on control variables vary depending on the treatment scenario:

Scenario	**MDA Constraints**	**EnvTre Constraints**
No Controls	$0 \leq u_C \leq 0$	$0 \leq u_W \leq 0$
MDA	$0 \leq u_C \leq 0.9$	$0 \leq u_W \leq 0$
EnvTre	$0 \leq u_C \leq 0$	$0 \leq u_W \leq 1$
MDA & EnvTre	$0 \leq u_C \leq 0.9$	$0 \leq u_W \leq 1$

Table A.2: Constraints on control variables in the different treatment scenarios considered in our analysis.

A.4.2 Numerical Methods

We use pseudospectral collocation to solve for the optimal dynamics of treatment and infection over time (see [31; 32; 33] for applications of this technique). Specifically, we approximate the optimal control model with a non-linear programming (NLP) problem, where we assume that our controls are approximated with an n^{th} degree polynomial over a period from 0 to T (the end of the planning horizon) [34]. The algorithm ensures that the residual error of the constraints is minimized at the collocation points. The collocation points and degree of polynomial are chosen to balance speed of convergence to a solution and numerical error; we used 60 collocation points. One advantage of this approach over more typical two-point boundary value methods, such as shooting, is that we can directly incorporate in the problem the constraints on the state and control variables [35]. This feature enables us to find optimal solutions that might reside on the boundary of the control set for a period of time. A second advantage is the ability to handle larger scale dynamical systems, such as the one in this paper with four states variables and two control variables. The solution method was implemented using TOMLAB (v. 8.4) [36; 37] and the accompanying PROPT toolbox [38]. The approximate NLP is solved using general-purpose nonlinear optimization packages KNITRO, SNOPT and NPSOL. We have included the source code that utilizes the TOMLAB/PROPT software at the end of the electronic supplementary material. Readers can download a trial version of the software to rerun our analysis.

A.5 Sensitivity Analyses

A.5.1 Basic Reproduction Ratio

According to Sokolow *et al.* [3], the expected range of the basic reproduction ratio R_0 for schistosomiasis ranges from 1 to 7. In the main results, the contact rates β_i and shedding rates λ_i were calibrated such that the basic reproduction ratio R_0 would be equal to 3.5. We investigate the range given in Sokolow *et al.* [3] by varying the contact rate from 1 infection per 700 water contacts ($R_0 = 1$), to 1 infection per 100 water contacts ($R_0 = 7$), and by varying the shedding rate from 1 intermediate host infection per 1944.4 sheds ($R_0 = 1$), to 1 intermediate host infection per 277.8 sheds ($R_0 = 7$). For the majority of the range of R_0 (at least between 2 and 7), our finding on the optimal substitution away from MDA due to the environmental treatment holds qualitatively (figure A.2, top two panels). For the environmental treatment, more time will be spent at the maximum level with a higher R_0, and more so if this higher R_0 is due to higher contact rates (figure A.2, bottom panel).

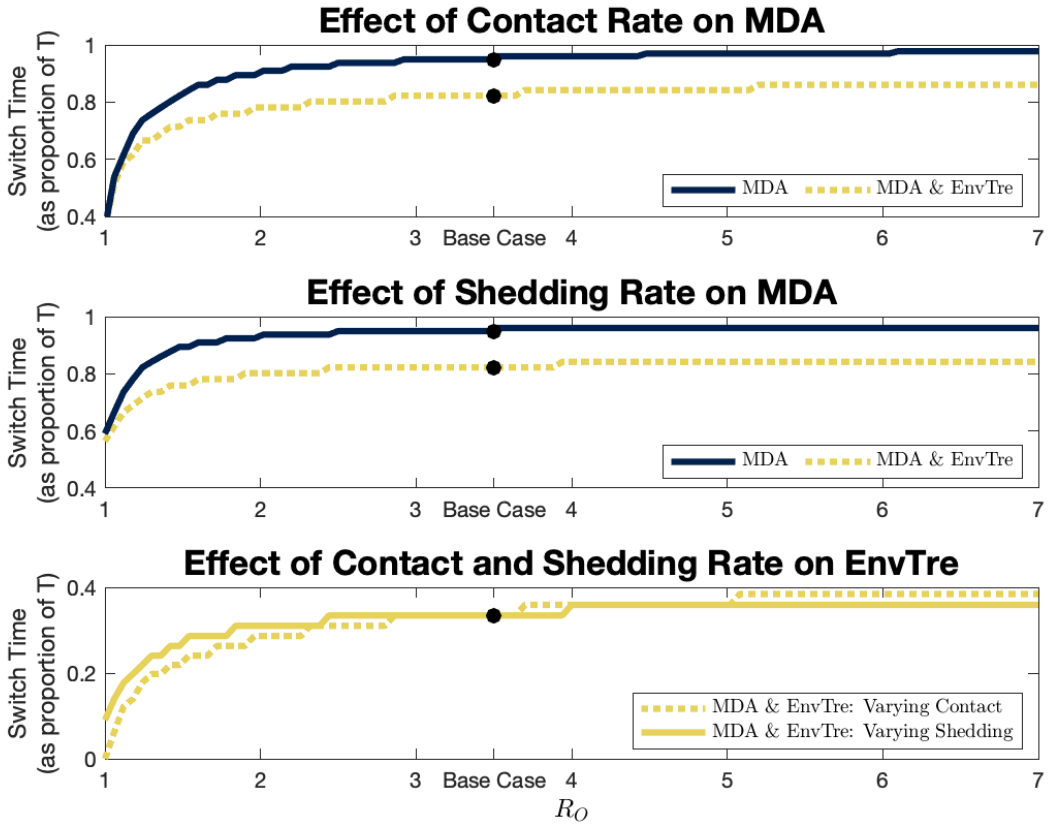

Figure A.2: Proportion of total treatment time spent at the maximum level of control (MDA or EnvTre) as a function of the basic reproduction ratio, R_0—which varies either with the contact rate (β_i) or the shedding rate (λ_i). The top panel shows the effect of the contact rate on children MDA, while the center panel shows the effect of the shedding rate on children MDA. The bottom panel shows the effect of the contact and shedding rate on the environmental treatment.

The contact rates (1 infection per 200 water contacts) and shedding rates (1 intermediate host infection per 555 sheds) in the main results are consistent with the literature [3; 19], however, there are many combinations of contact and shedding rates that could yield a R_0 of 3.5, even if we keep the biological parameters identical across children and adults. Our sensitivity analyses suggest that the substitution away from MDA due to EnvTre remains the same regardless of how the contact rate is relative to the shedding rate (figure A.3, top panel). Contact rates play a bigger role than shedding rates in determining the amount of time spent at maximum treatment (figure A.3). As the contact rate (β_i) increase relative to the shedding rate (λ_i), and keeping R_0 constant, more time will be spent utilizing both MDA and EnvTre controls.

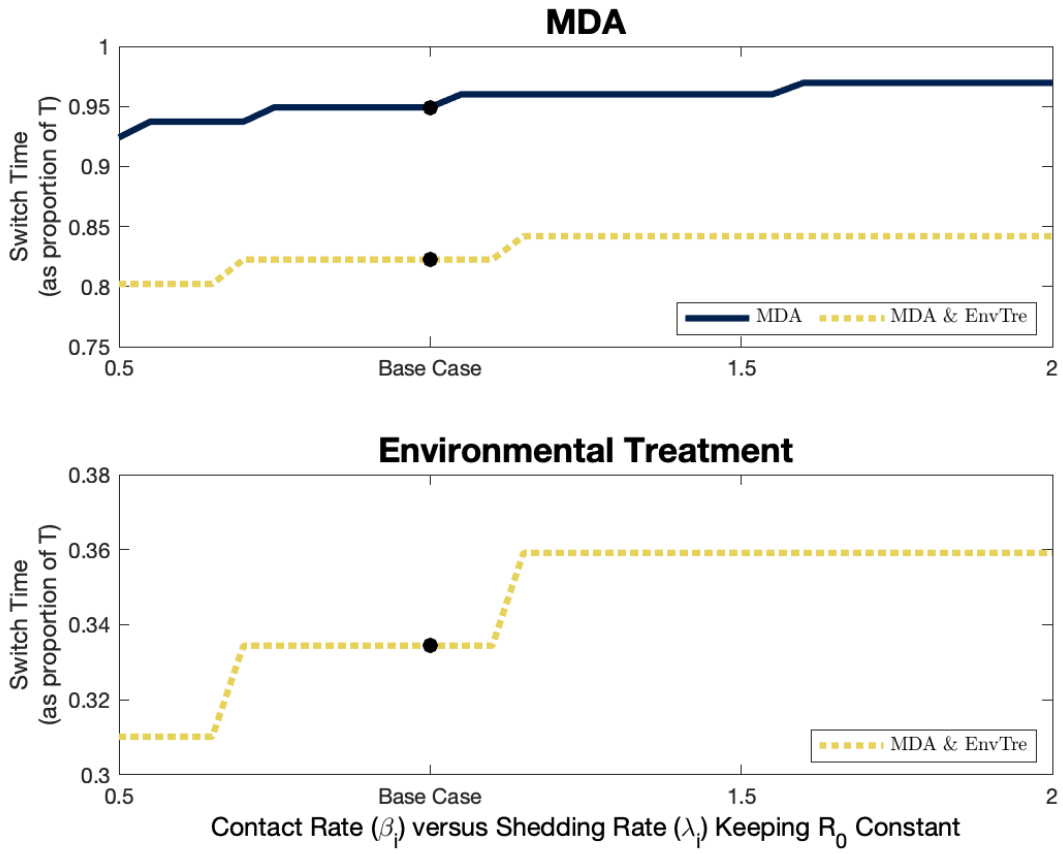

Figure A.3: This figure shows the proportion of total time spent at the maximum level of MDA (i.e., $u_C = 0.9$, top panel) and the proportion of total time spent at the maximum level of EnvTre (i.e., $u_W = 1$, bottom panel) as a function of the contact rate relative to the shedding rate, keeping R_0 constant at 3.5. A value of 0.5 on the x-axis represents a case where the contact rate is half of its base case value, while the shedding rate is twice its value; a value of 2 on the x-axis is the converse.

A.5.2 Proportion of Children in Total Population

In our simulation, the proportions of children versus adults are based on data from Senegal where 40% of the population are children (0-14 years) and 60% are adults (15 years and over). As the proportion of children in total population increases, the damage parameter associated with infectious children (c_C) increases while the one associated with infectious adults (c_A) decreases because of the way they are constructed (i.e., a constant times the respective proportion), however MDA cost parameter associated with children (c_{PC}) increases proportionally. As a result, there is no change in the amount of time spent treating children (figure A.4, top panel). Furthermore, only when the proportion of children becomes low enough, and in turn the expenditures associated with child MDA, we find a discrete shift up in the amount of time spent treating the environment (figure A.4, bottom panel).

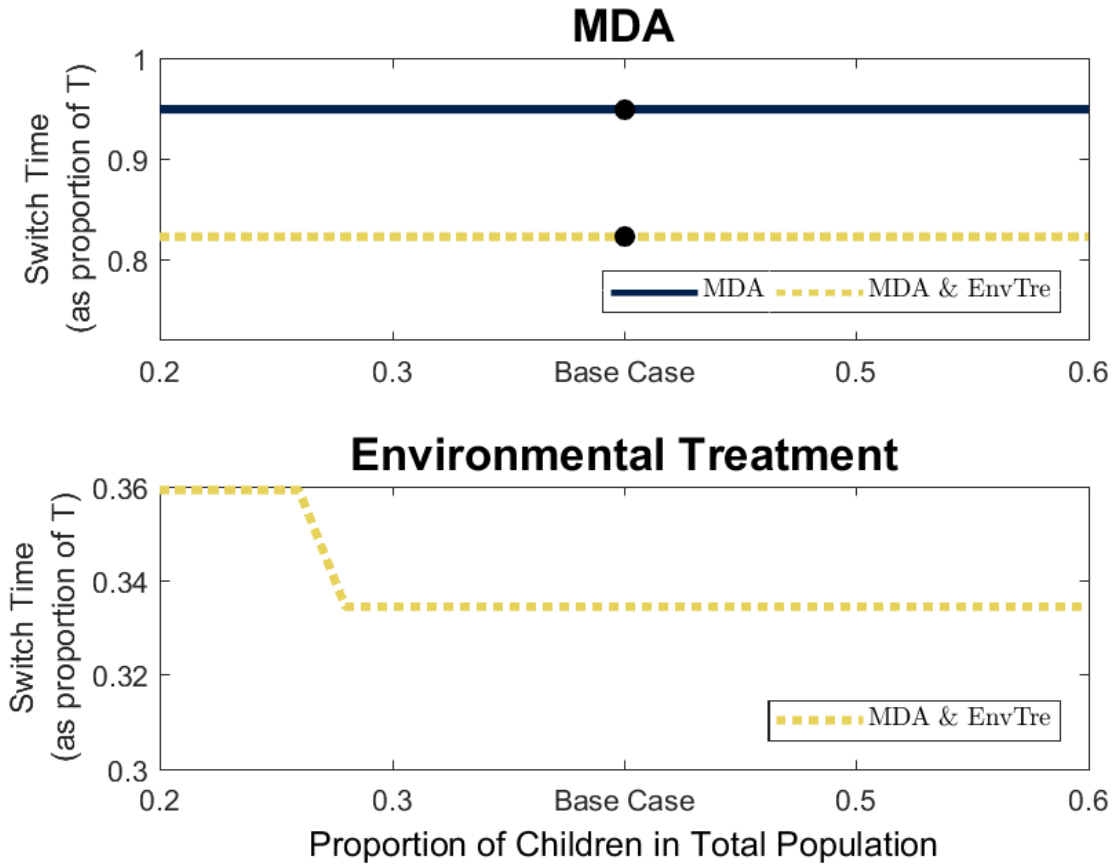

Figure A.4: This figure shows the proportion of total time spent at the maximum level of MDA (i.e., $u_C = 0.9$, top panel) and the proportion of total time spent at the maximum level of EnvTre (i.e., $u_W = 1$, bottom panel) as a function of the proportion of children in the total population.

A.5.3 Initial Conditions

In our analysis, all state variables are at their pre-treatment, long-term, levels. By varying the initial conditions from 10% to 100% of these levels, we find that the optimal substitution away from MDA to EnvTre remains approximately the same, except when values drop below the 20% threshold (figure A.5, top panel). Since moving further away from the long-term value implies a lower number of infected intermediate hosts in the environment, the switch time of EnvTre will reduce the further we are from the steady-state values (figure A.5, bottom panel).

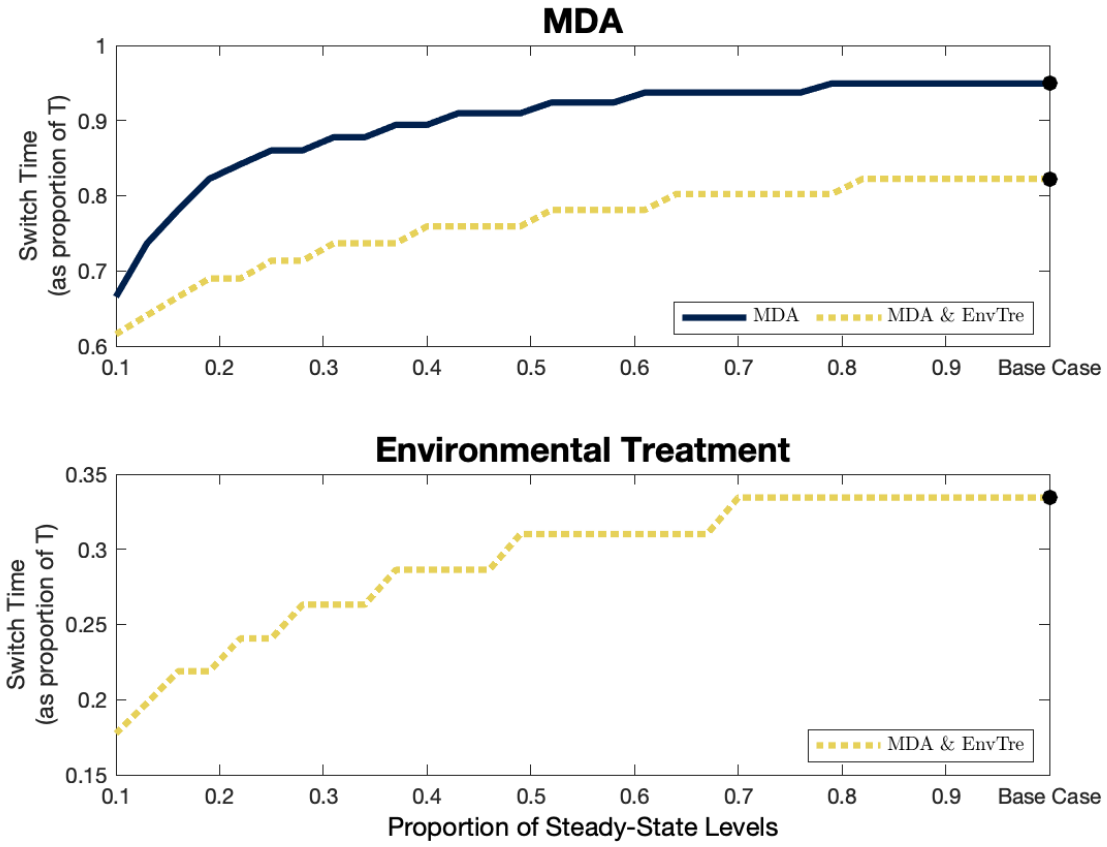

Figure A.5: This figure shows the proportion of total time spent at the maximum level of MDA (i.e. $u_C = 0.9$, top panel) and the proportion of total time spent at the maximum level of EnvTre (i.e. $u_W = 1$, bottom panel) as a function of the initial conditions. Initial conditions range from 10% to 100% of the long-term values for state variables I_C, I_A, and W; the snail population size $(X + W)$ is assumed to be at its carrying capacity in all simulations.

A.5.4 Discount Rate

The higher the discount rate r the more weight is placed on damages, treatment, and transportation and management costs incurred early in the program. By varying the discount rate, we find that it does not vary the optimal substitution away from MDA due to EnvTre in the $r = [0, 0.2]$ range (base case, $r = 0.04$) (figure A.6, top panel) Conversely, the level of EnvTre reduces as the discount rate increases indicating that EnvTre has long-term benefits that are attenuated when the discount rate increases (figure A.6, bottom panel). For a discussion about discount rates for health outcomes, see [39] and [40].

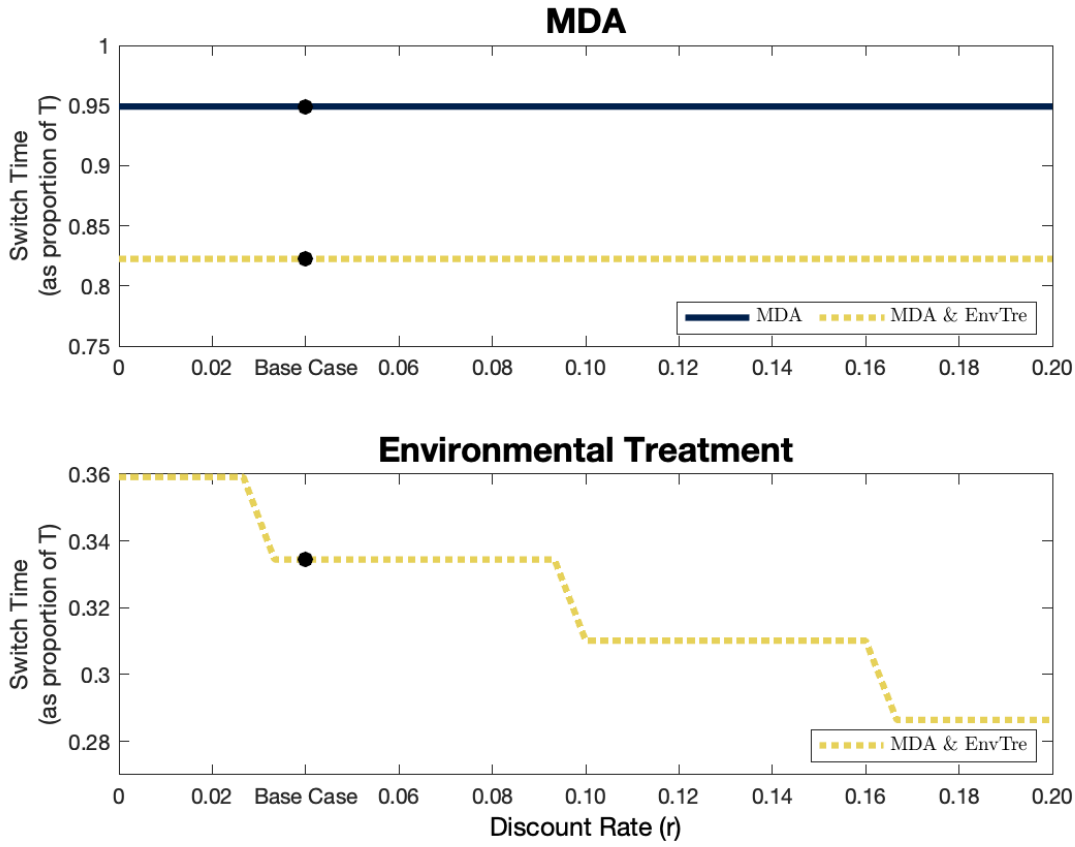

Figure A.6: This figure shows the proportion of total time spent at the maximum level of MDA (i.e. $u_C = 0.9$, top panel) and the proportion of total time spent at the maximum level of EnvTre (i.e. $u_W = 1$, bottom panel) as a function of the discount rate.

A.5.5 Value of a DALY

The value of a DALY was set to be approximately the median GDP per capita of an African country. Intuitively, the switch time for both MDA (figure A.7, top panel) and EnvTre (figure A.7, bottom panel) will increase as the value of a DALY increases. Because, everything else equal, damages are relatively more important than treatment costs if the value of a DALY increases, more treatment will occur. The substitution from MDA to EnvTre will decrease as the value of a DALY increases since the effect of EnvTre is only indirect compared to MDA that directly, and immediately, reduces disease burden.

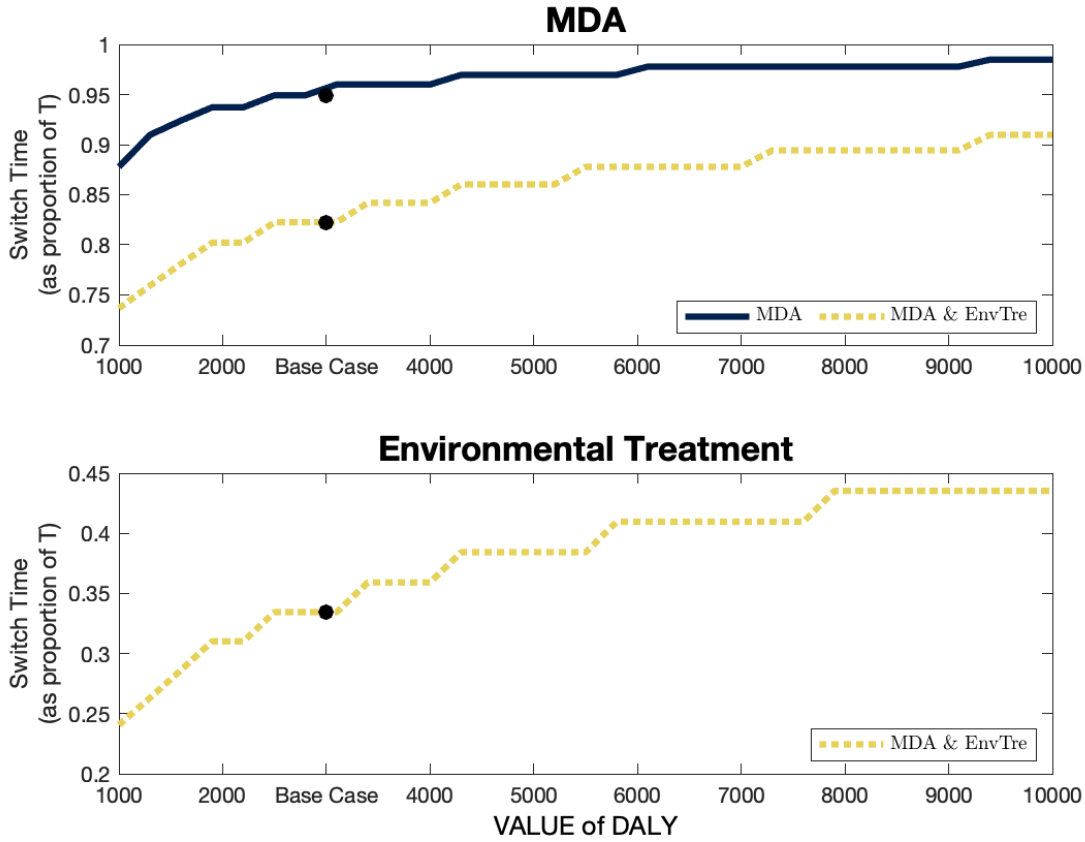

Figure A.7: This figure shows the proportion of total time spent at the maximum level of MDA (i.e. $u_C = 0.9$, top panel) and the proportion of total time spent at the maximum level of EnvTre (i.e. $u_W = 1$, bottom panel) as a function of the value of DALYs.

A.5.6 Effectiveness of Controls

MDA and EnvTre have a limited effectiveness. In these sensitivity analyses, we investigate how our results would be affected by a worsened or improved effectiveness. We find that if the effectiveness of MDA improves, the substitution away from MDA due to EnvTre would increase (figure A.8, top panel) since a higher MDA effectiveness would reduce the optimal amount of time needed to reduce disease burden in humans. Though one could expect an improved effectiveness to have the same effect on both the MDA and MDA & EnvTre cases, the effect is more pronounced in the latter because the reinfection of humans is lower when EnvTre is utilized. Effectiveness of MDA has no impact on the optimal amount of EnvTre (figure A.8, bottom panel). As for the effectiveness of EnvTre, it has no effect on the substitution away from MDA due to EnvTre (figure A.9, top panel), and an improved effectiveness in EnvTre control would reduce the optimal amount of EnvTre treatment (figure A.9, bottom panel).

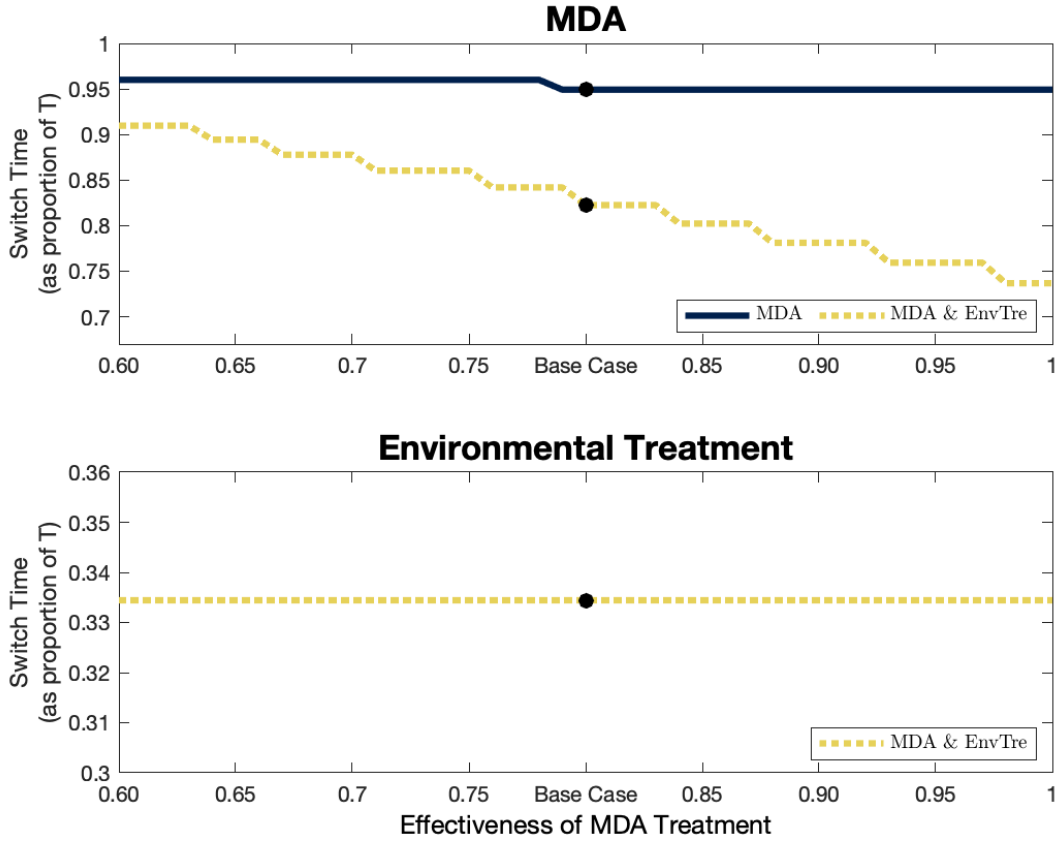

Figure A.8: This figure shows the proportion of total time spent at the maximum level of MDA (i.e. $u_C = 0.9$, top panel) and the proportion of total time spent at the maximum level of EnvTre (i.e. $u_W = 1$, bottom panel) as a function of the effectiveness of MDA control.

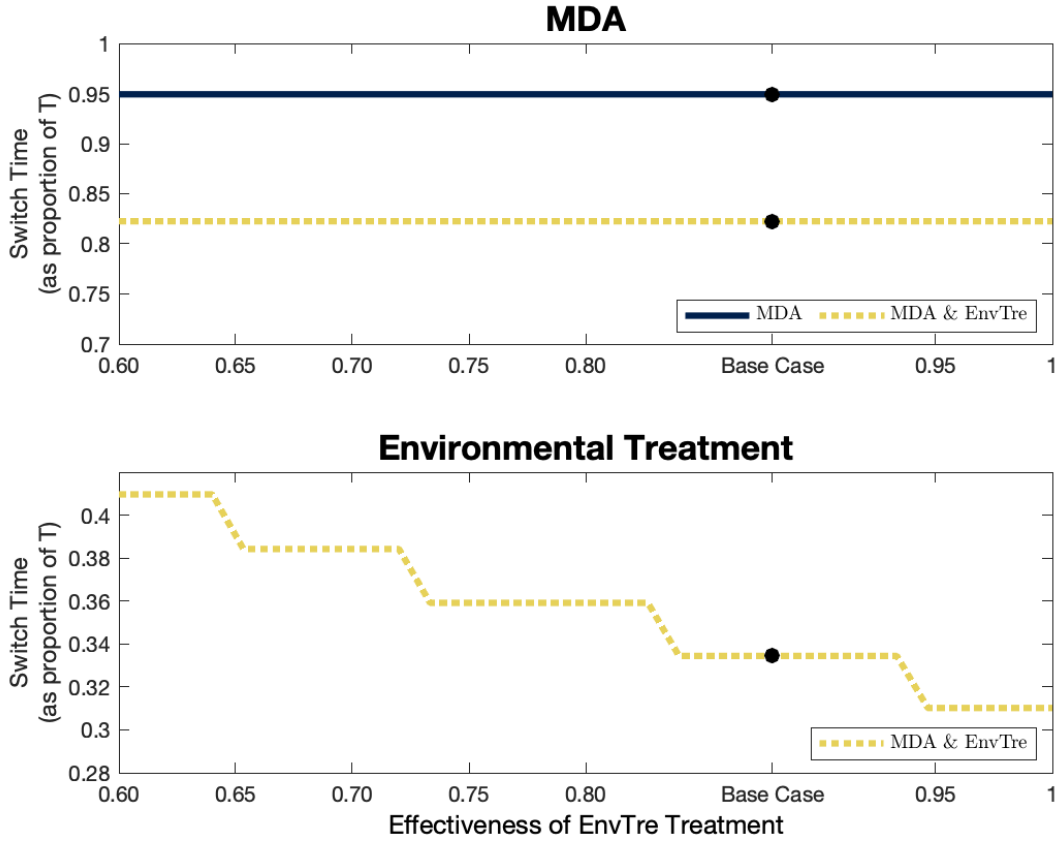

Figure A.9: This figure shows the proportion of total time spent at the maximum level of MDA (i.e. $u_C = 0.9$, top panel) and the proportion of total time spent at the maximum level of EnvTre (i.e. $u_W = 1$, bottom panel) as a function of the effectiveness of EnvTre.

A.5.7 Growth Rate of Intermediate Hosts

We find that a variation of 50 percent around our base growth rate of intermediate hosts does not change our results (figure A.10). The values considered in our sensitivity analysis did not vary our results because the growth rate remains several orders of magnitude higher than the natural death rate of snails. Only when the growth rate is of the same order of magnitude as the death rate do we find measurable changes; when growth rate drops below the death rate, it follows that treatment will reduce drastically.

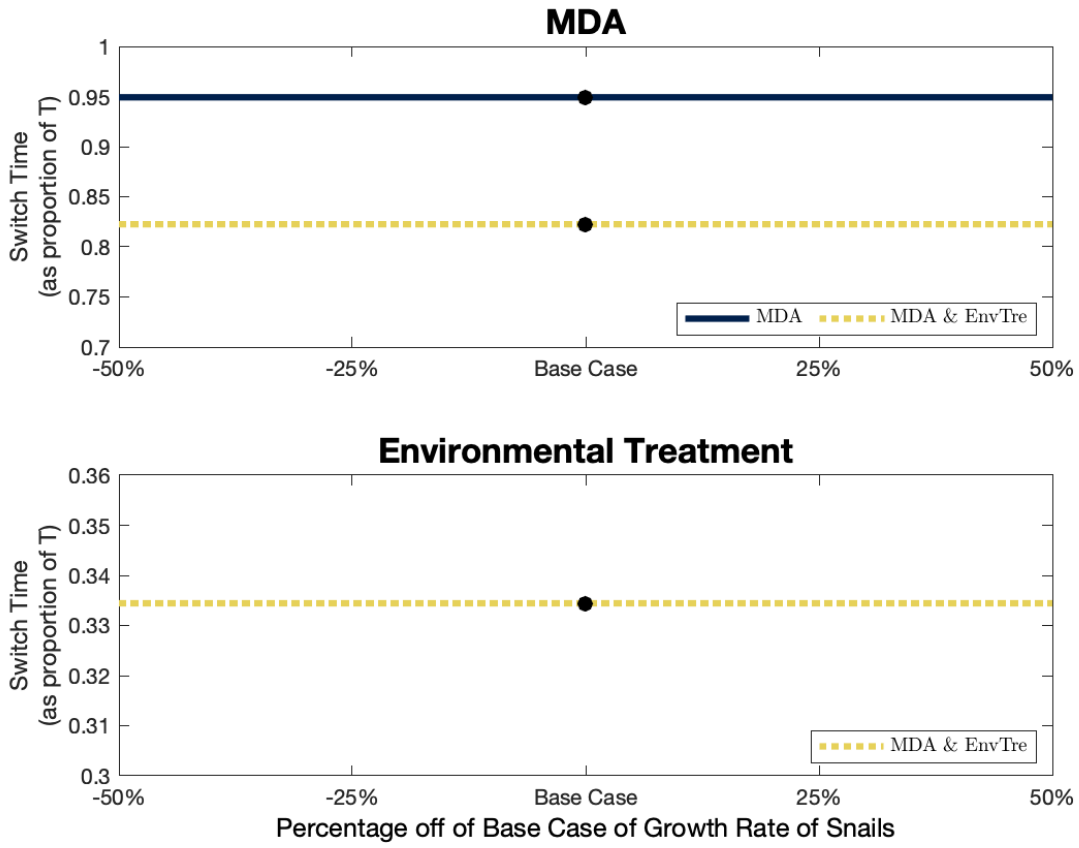

Figure A.10: This figure shows the proportion of total time spent at the maximum level of MDA (i.e. $u_C = 0.9$, top panel) and the proportion of total time spent at the maximum level of EnvTre (i.e. $u_W = 1$, bottom panel) as a function of the growth rate of snails.

A.5.8 Time Horizon

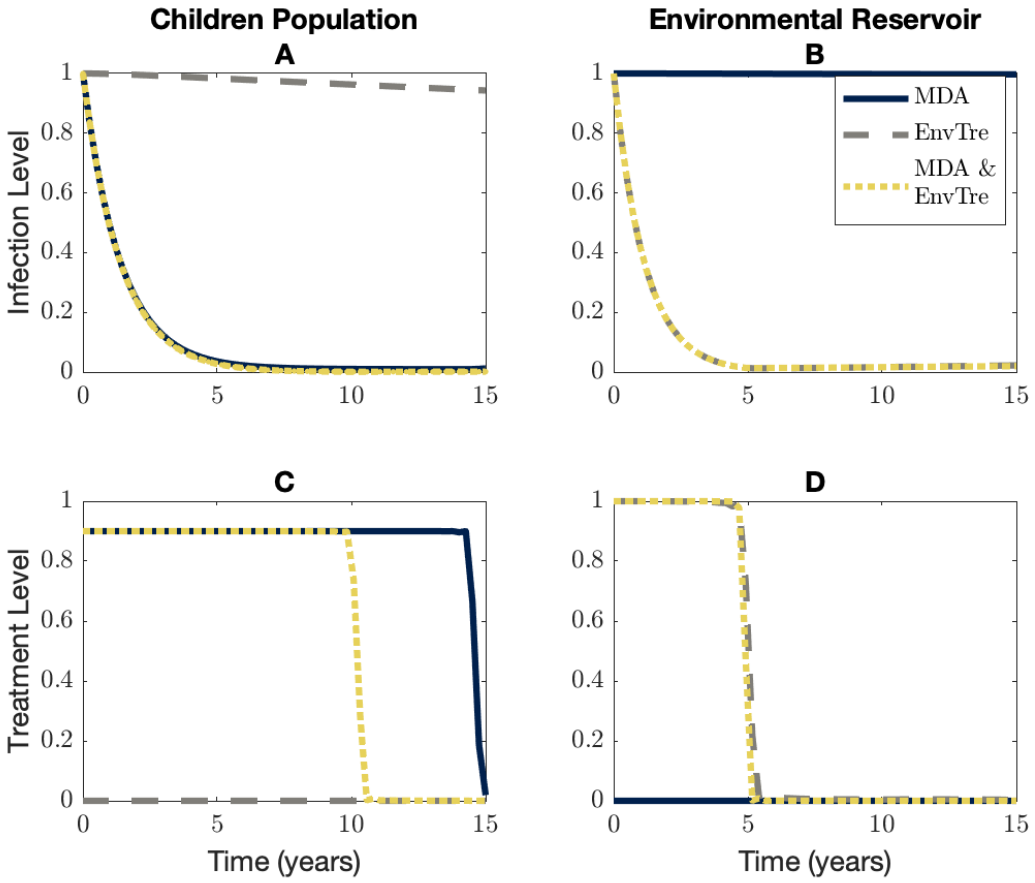

Figure A.11: This figure shows the change over time of the infection prevalence in the child population (panel A), the infection prevalence in the intermediate host population (panel B), the optimal path of drug treatment (panel C) and the optimal path of environmental treatment (panel D) for when $T = 15$ years. Infection prevalence is expressed as a proportion of its respective pre-treatment, long-term, value.

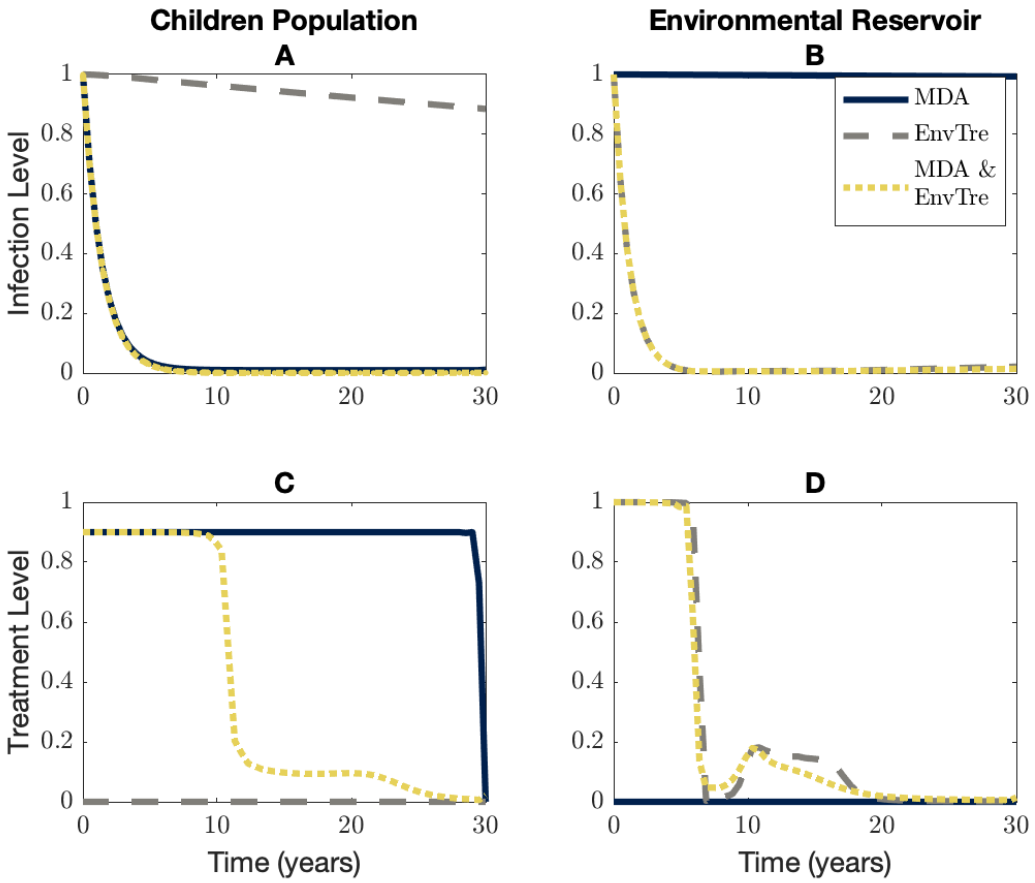

Figure A.12: This figure shows the change over time of the infection prevalence in the child population (panel A), the infection prevalence in the intermediate host population (panel B), the optimal path of drug treatment (panel C) and the optimal path of environmental treatment (panel D) for when $T = 30$ years. Infection prevalence is expressed as a proportion of its respective pre-treatment, long-term, value.

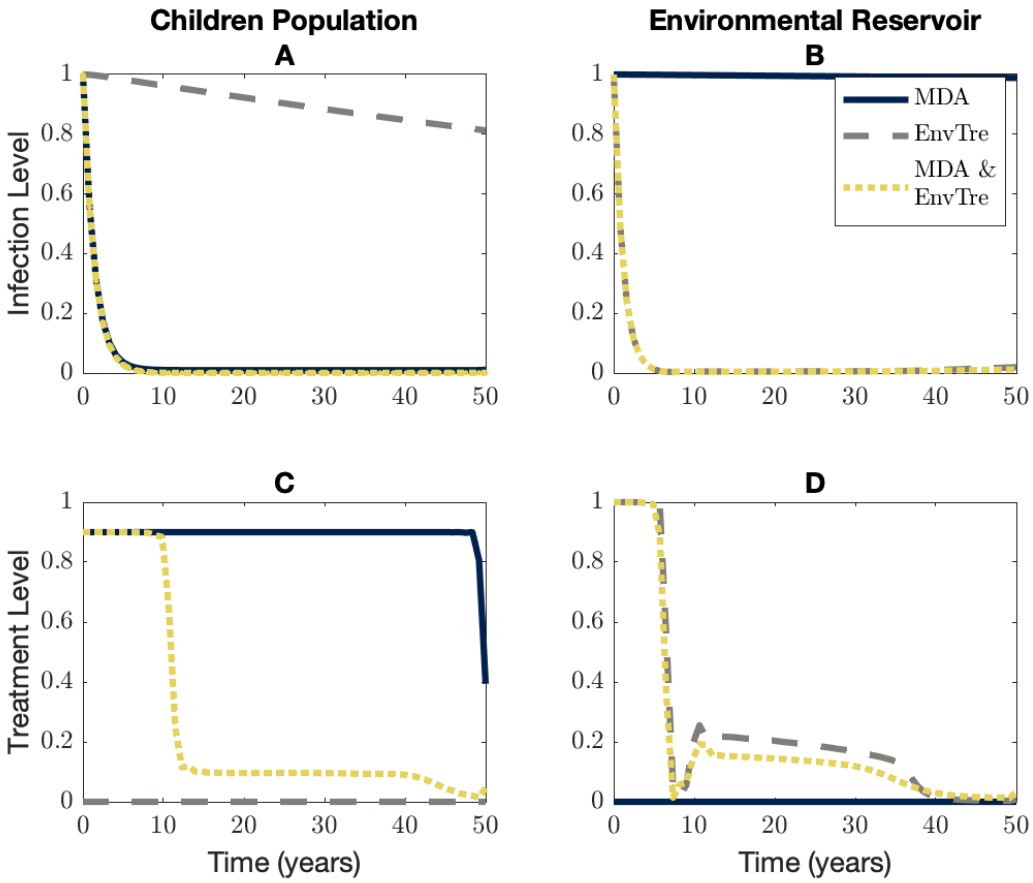

Figure A.13: This figure shows the change over time of the infection prevalence in the child population (panel A), the infection prevalence in the intermediate host population (panel B), the optimal path of drug treatment (panel C) and the optimal path of environmental treatment (panel D) for when $T = 50$ years. Infection prevalence is expressed as a proportion of its respective pre-treatment, long-term, value.

B Chapter 2

B.1 Parameterization

B.1.1 Epidemiological Model

According to Diekmann et al. [16], 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-jurisdiction SEIR model in this paper—one needs to use the next-generation matrix and find its dominant eigenvalue to find the R_0 [16]. 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 [7]; the next-generation matrix is equal to $-FV^{-1}$. In the model presented in this paper,

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

where the four rows of F and V refer to the E_1, I_1, E_2 and I_2 equations, respectively. Note that both matrices F and V are derived under the assumption of introducing a single exposed individuals in an otherwise susceptible population (for more details on how to construct the next-generation matrix in a SEIR model, see [41]). Given we assume that $\beta_{11} = \beta_{22} = \beta_{ii}$ and $\beta_{12} = \beta_{21} = \beta_{ij}$, and when in our main analysis we let $\varphi_1 = \varphi_2 = \varphi$, the basic reproduction ratio of our model simplifies to,

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

for $i = 1, 2$, $j = 1, 2$, and $i \neq j$. We set the basic reproduction ratio $R_0 = 1.43$, according to estimates of the R_0 from Li et al. [42] and using estimates of the effect of nonpharmaceutical interventions on the R_0 from Tian et al. [43]. We assume a mean recovery period ($\frac{1}{\gamma}$) of 5 days [44], and a case-fatality ratio of 1.78% (adjusted for misreporting, see [45]) to calibrate the rate of disease induced mortality, φ. Parameters β_{ii} and β_{ij} are then calibrated assuming what Tian et al. [43] 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 [46]); 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 [44].

B.1.2 Economic Model

To quantify damages, we use the value of statistical life recommended by the Environmental Protection Agency.[20] The disability weight[21] 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).[22] This disability weight thus allows for a comparison between the individuals that are infected with the disease but do not die, and the individuals that die from its complications.

Expenditures related to the pharmaceutical intervention are based off estimates of vaccine costs. Numerous governments around the world, including the U.S. federal government, have contracted biotech companies producing COVID-19 vaccines; governments pay money in exchange of a guaranteed number of doses of COVID-19 vaccines. These estimates and the prices of current influenza vaccine turns out to be approximately 20

[20] See "What value of statistical life does EPA use?" from the U.S. Environmental Protection Agency [47].

[21] 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/`.

[22] For more details on how COVID-19's disability resembles lower respiratory tract infections, see [48].

U.S. dollars per dose, with two doses per individual; this is the value we chose in our analysis.[23]

The value of the workability cost[24] 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 [40] for a discussion about discounting health-related expenditures).

B.1.3 Parameter Levels

Table B.1 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}).[25]
β_{ij}	4.36	Transmission rate across states (month^{-1}).[25]
σ	10.14	Rate at which infected individuals become infectious (month^{-1}).[26]
γ	6.08	Rate of recovery (month^{-1}).[26]
ω	0.17	Rate at which immunity is lost (month^{-1}).[27]
φ	0.11	Rate of disease induced mortality (month^{-1}).[28]
w	0.13	Disability weight associated with the disease (unitless).[29]
q_V	0.65	Efficiency of vaccines (proportion).[30]
r	0.0013	Discount rate (month^{-1}).[31]
c_V	40	Cost of treating one individual via vaccine (US Dollars).[32]
c_A	10×10^3	Workability cost (US Dollars).[33]
c	10×10^6	Value of statistical life (US Dollars).[34]

Table B.1: Parameter levels used in the numerical simulation.

[23]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`.

[24]Inspired by the paper of Ryan et al. [49] where the authors show the implications of policy adjustment costs for fisheries management

B.2 Optimization

B.2.1 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 [43] (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:

[25]Calibrated using a R_0 estimate from Li et al. [42] and estimates of effects of nonpharmaceutical interventions from Tian et al. [43]; 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.

[26]Using estimates from Davies et al. [44]; this represents a 3-day latency period and a 5-day recovery period.

[27]Representing a 6-month immunity period in the scenarios where we assume immunity is not permanent; based on [50].

[28]Calibrated by using a case-fatility rate of 1.78%. Adjusted for mis- and under-reporting; see [45].

[29]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 [48].

[30]Following Buckner et al. [51], we base this parameter value on the efficiency of the influenza vaccine [52]. Note that the Department of Health and Human Services, Food and Drug Administration (FDA), and Center for Biologics Evaluation and Research (CBER) [53] requires that a future COVID-19 vaccine must have an effectiveness of at least 50%.

[31]Based on results from John et al. [40] 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 $r = 0.0013$.

[32]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`.

[33]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.

[34]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 [47] uses: approximately $7.4 million ($2006) which is equivalent to approximately $9.54 million ($2020).

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

Table B.2: 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:

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

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

B.2.2 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:

$$0 \leq S_i \leq N_i \leq 1 \text{ for } i = 1, 2; \tag{B.2a}$$

$$0 \leq E_i \leq N_i \leq 1 \text{ for } i = 1, 2; \tag{B.2b}$$

$$0 \leq I_i \leq N_i \leq 1 \text{ for } i = 1, 2; \tag{B.2c}$$

$$0 \leq R_i \leq N_i \leq 1 \text{ for } i = 1, 2; \tag{B.2d}$$

$$S_i + E_i + I_i + R_i = N_i \leq 1 \text{ for } i = 1, 2. \tag{B.2e}$$

Control variables are modelled as direct controls (see examples in [54; 55; 56]) and can be interpreted as a reduction in the number of susceptible individual in a given time period (i.e. a month). Formally, the constraints on the control variables are given by:

$$0 \leq u_{V_i} \leq S_i \text{ for } i = 1, 2. \tag{B.3}$$

Because of a limited supply of vaccines (see details below), the physical upper-bound on constraints (B.3) will only be binding when capacity constraint is nonbinding. When this occurs, it means that there are fewer susceptible individuals than there are available vaccines.

B.2.3 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 before the supply has had a chance to ramp up. Hence, the control variables are also subject to

$$u_{V_1} + u_{V_2} \leq \bar{u}_V; \tag{B.4}$$

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

$$u_{V_i} \leq \left(\frac{N_i}{N_1 + N_2} \right) \bar{u}_V \quad \text{for } i = 1, 2. \tag{B.5}$$

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

B.2.4 Numerical Methods

Pseudospectral collocation approximates the continuous time optimal control model with a constrained nonlinear programming problem (see [31; 32; 57; 58] for other applications of this technique). The dynamic controls to our problem—i.e. the vaccine allocation—are approximated by a polynomial of degree n (determined by the number of collocation points) over a period from $t = 0$ (date at which the vaccine starts to be administered) to $t = T$ (assumed to be four months after the vaccine administration) [34]. The residual error of the constraints is minimized by the algorithm at the n collocation points, where n is chosen to have a reasonable speed of convergence to a solution and a low numerical

error. Here, we chose 60 collocations points. In this sort of problem, the main advantage of this approach over more usual methods to solve such two-point boundary problems, such as shooting methods, is that nonnegativity constraints (e.g. on the number of infected individuals) and upper-bound constraints (mimicking e.g. vaccine capacity constraints) on state and control variables can be directly incorporated in the problem [35]. This method thus allows us to find optimal solutions that may lay on the boundary of the control set for a certain period of time. For COVID-19 vaccines, this is likely due to the scarcity of the supply of vaccine in the short-term. Another advantage of this method is the ability to deal with large-scale dynamical systems, such as the one presented here with ten state variables and two control variables. The solution was found using TOMLAB (v. 8.4) [36; 37] and the accompanying PROPT toolbox [38]. The approximate nonlinear programming problem is solved using general-purpose nonlinear optimization packages (e.g. KNITRO, SNOPT and NPSOL).

B.3 Figures: Homogeneous Demographic Characteristics

B.3.1 Compliance and Noncompliance to the Travel Restrictions

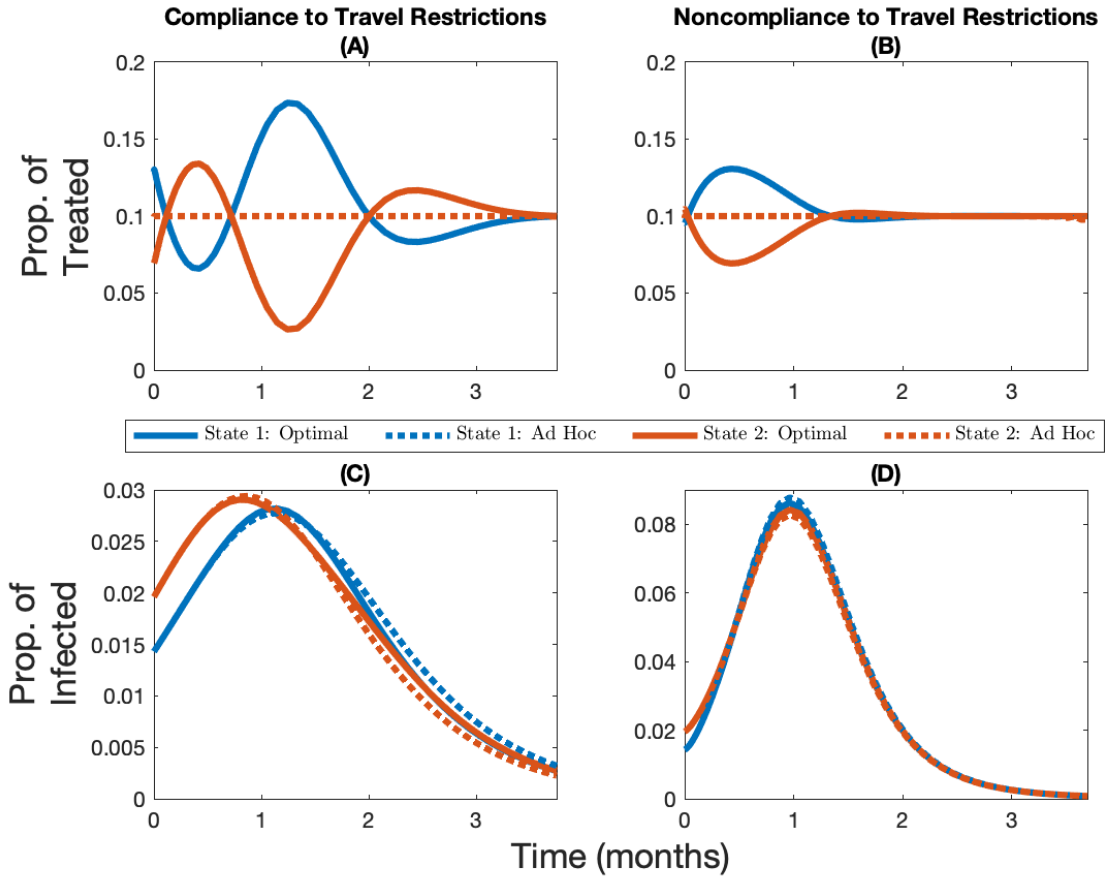

Figure B.1: **Permanent immunity with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the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

B.3.2 Vaccine Capacity Constraints when Immunity is Permanent

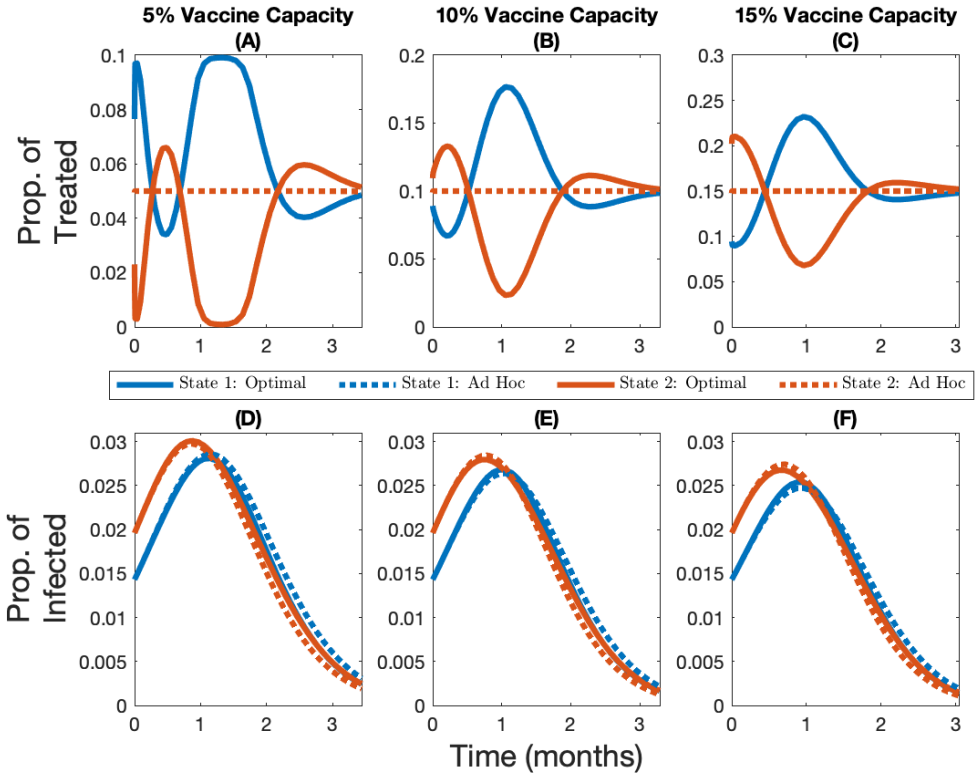

Figure B.2: **Permanent immunity and compliance to travel restrictions with 5%, 10%, and 15% vaccine capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened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

B.3.3 Vaccine Capacity Constraints when Immunity is Temporary

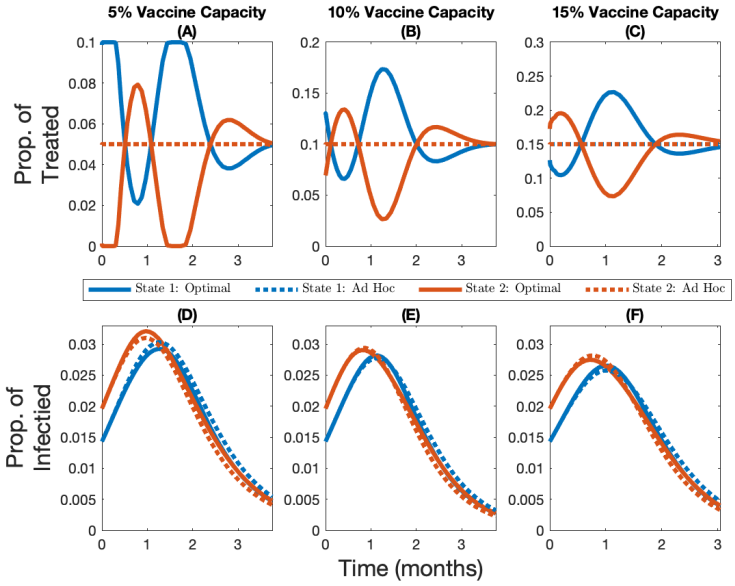

Figure B.3: **Temporary immunity and compliance to travel restrictions with 5%, 10%, and 15% vaccine capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether capacity is 5% (panels A and D), 10% (panels B and E), or 15% (panels C and F), for the case where immunity lasts six months and there is compliance to travel restrictions.

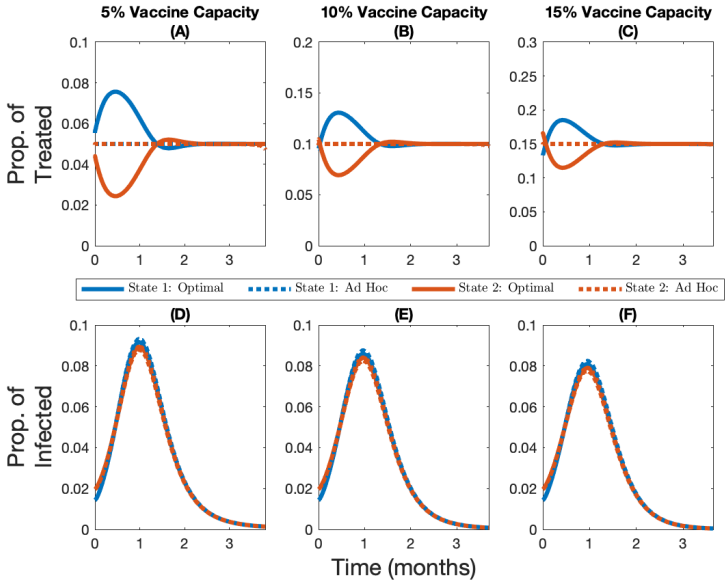

Figure B.4: **Temporary immunity and noncompliance to travel restrictions with 5%, 10%, and 15% vaccine capacity.** Change over time in the optimal and *ad hoc* allocations (panels A, B, and C) and the corresponding infection levels (panels D, E, and F) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether capacity is 5% (panels A and D), 10% (panels B and E), 15% (panels C and F), for the case where immunity lasts six months and there is no compliance to travel restrictions.

B.3.4 Permanent vs Temporary Immunity

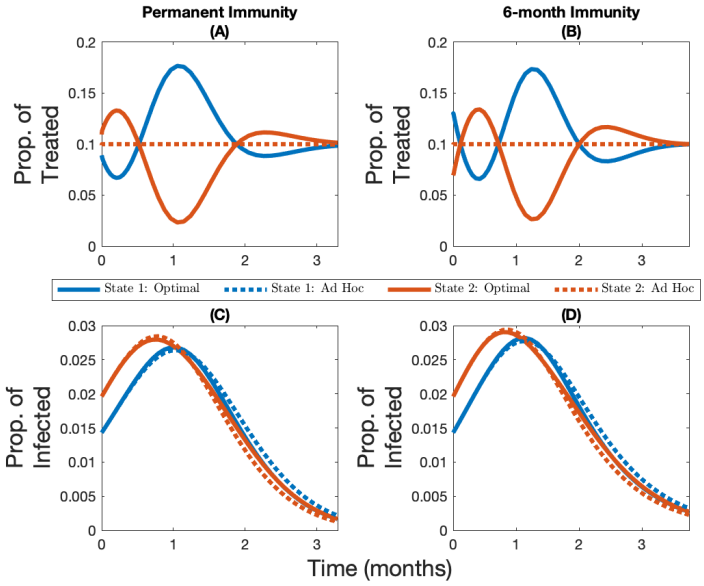

Figure B.5: **Compliance to travel restrictions with permanent and temporary immunity.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the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

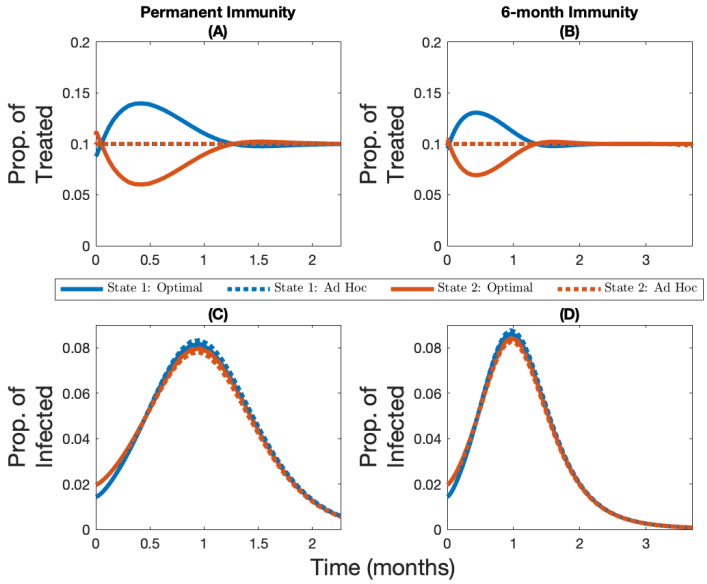

Figure B.6: **Noncompliance to travel restrictions with permanent and temporary immunity.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether immunity is permanent (panels A and C) or lasts six months (panels B and D) for the case where the vaccine capacity constraint is 10% and there is no compliance to travel restrictions.

B.3.5 Cumulative Infection Levels

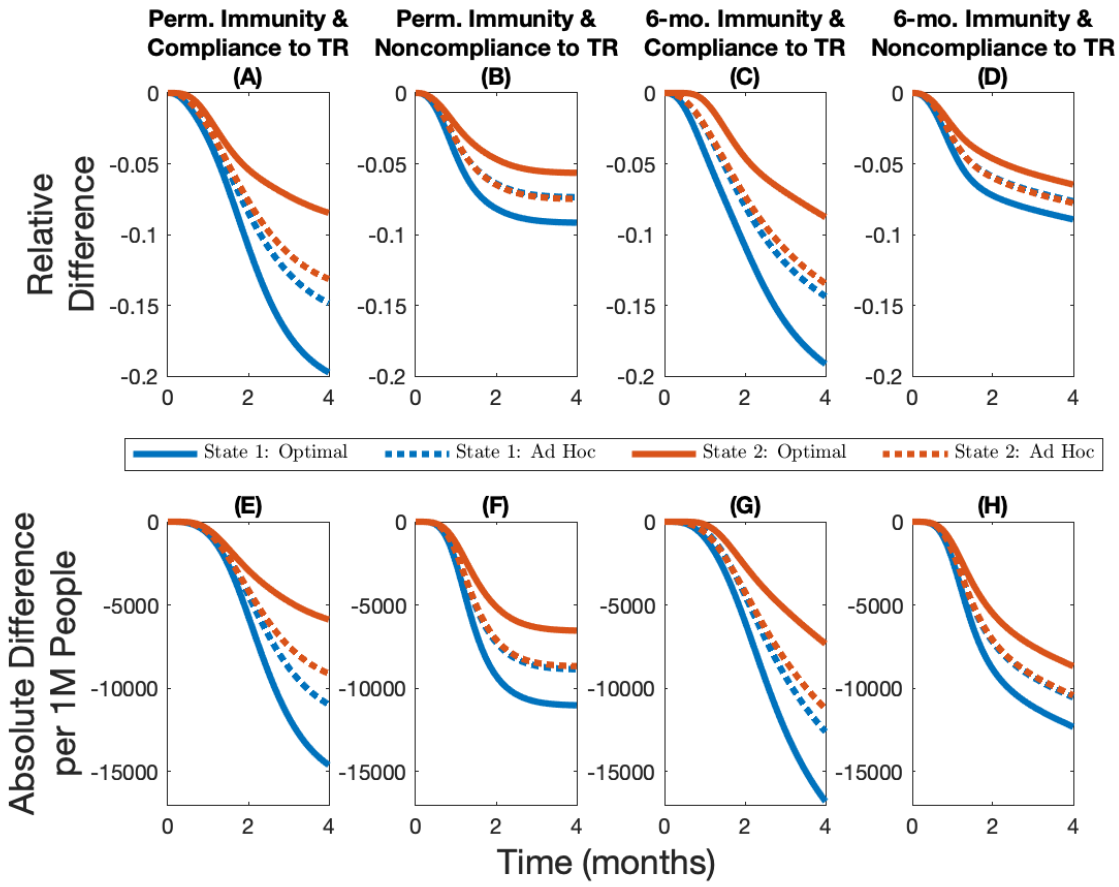

Figure B.7: **Epidemiological outcomes under different scenarios with a low vaccine supply.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-vaccine case for different immunity–travel restrictions scenarios and for when vaccine capacity is 5%.

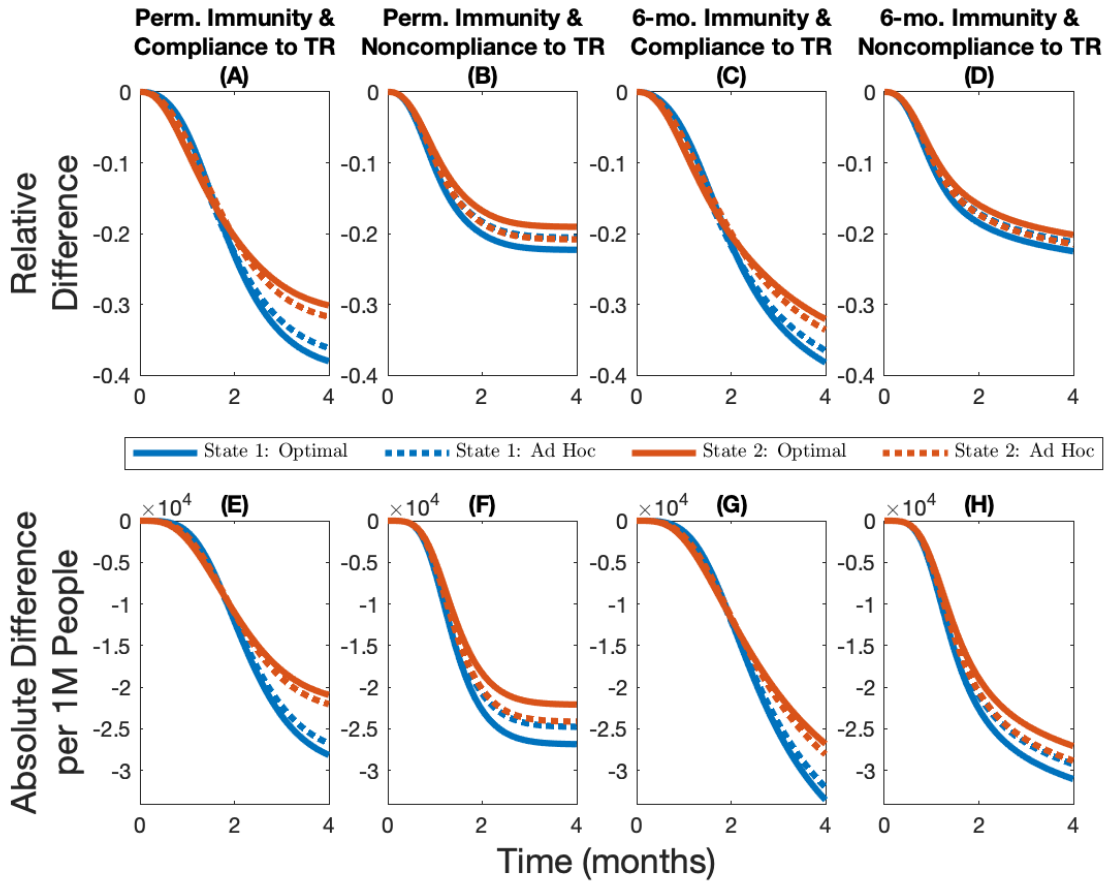

Figure B.8: **Epidemiological outcomes under different scenarios with a high vaccine supply.** Cumulative relative difference (panels A, B, C, and D) and cumulative absolute difference per 1M people (panels E, F, G, and H) between the number of infections in different allocations rules and the no-vaccine case for different immunity–travel restrictions scenarios and for when vaccine capacity is 15%.

B.4 Figures: Heterogeneous Demographic Characteristics

B.4.1 Heterogeneous Case-Fatality Ratio

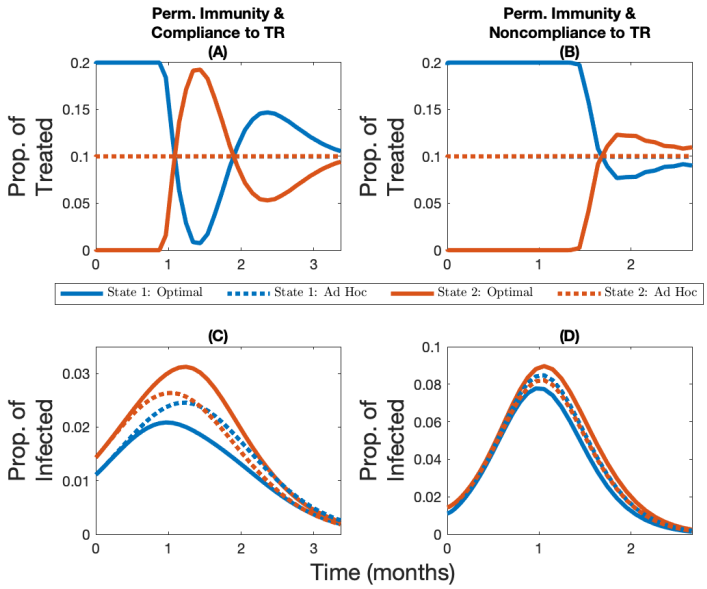

Figure B.9: **Vaccine allocation with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether there is compliance to travel restrictions (panels A and C) or not (panels B and D) for the case where the vaccine capacity constraint is 10%, immunity is permanent, and where the heterogeneity in the system comes from a varying case-fatality ratio (State 1 has a case-fatality ratio 1% higher than State 2).

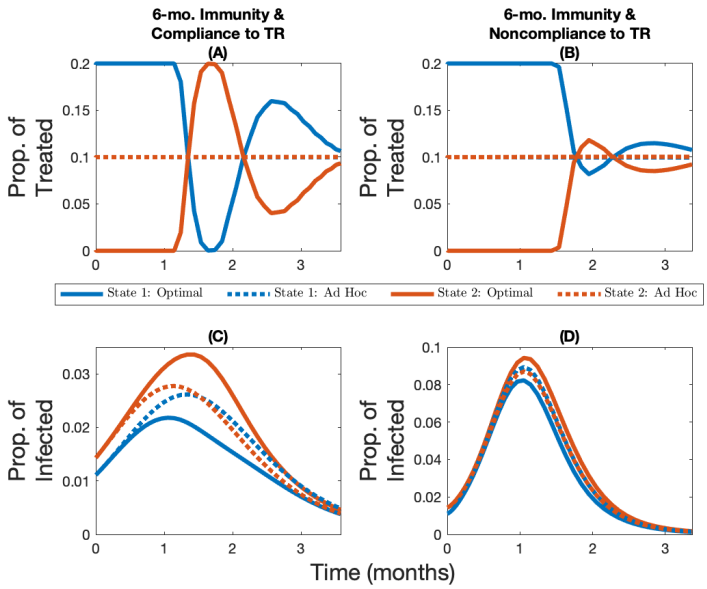

Figure B.10: **Vaccine allocation with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the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

B.4.2 Heterogeneous Contact Rate

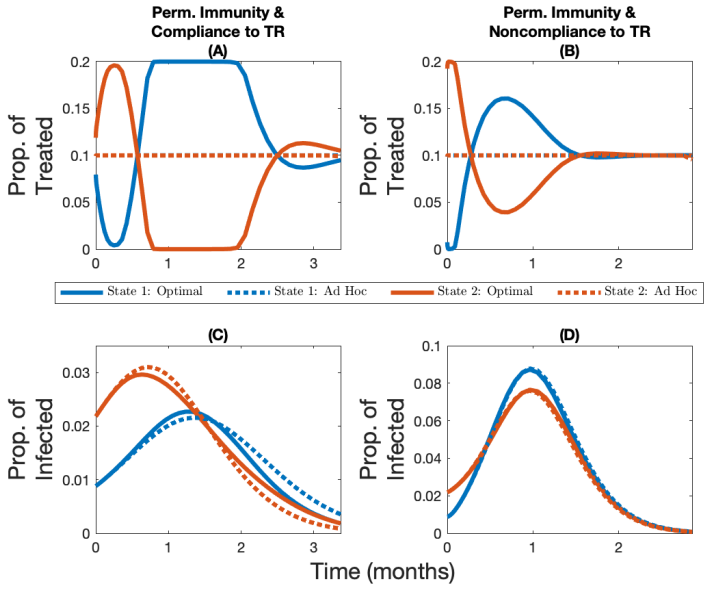

Figure B.11: **Vaccine allocation with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether there is compliance to travel restrictions (panels A and C) or not (panels B and D) for the case where the vaccine capacity constraint is 10%, immunity is permanent, and where the heterogeneity in the system comes from a varying contact rate (State 2 has a higher contact rate than State 1).

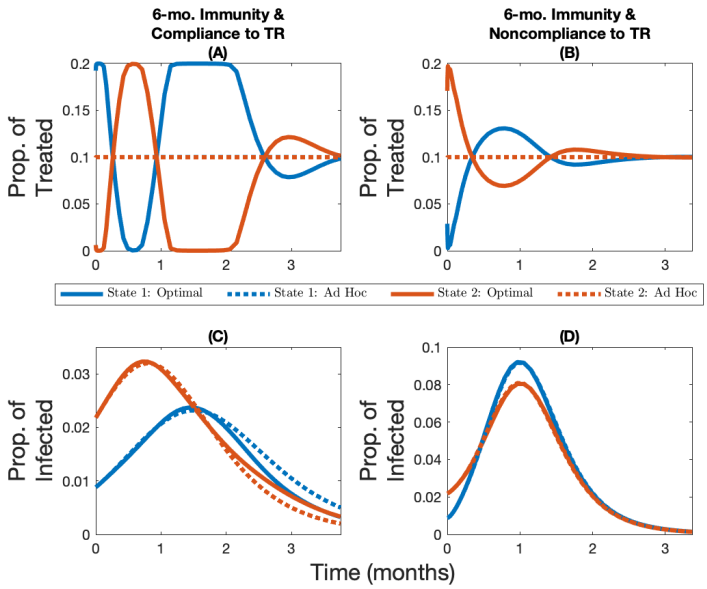

Figure B.12: **Vaccine allocation with and without compliance to travel restrictions.** Change over time in the optimal and *ad hoc* allocations (panels A and B) and the corresponding infection levels (panels C and D) for State 1 (in blue, the initially lowest-burdened state) and State 2 (in red, the initially highest-burdened state) depending on whether there is compliance to travel restrictions (panels A and C) or not (panels B and D) for the case where the vaccine capacity constraint is 10%, immunity lasts six months, and where the heterogeneity in the system comes from a varying contact rate (State 2 has a higher contact rate than State 1).

B.5 Figures: Robustness of Spatial allocations

B.5.1 Base Case: Homogeneous Demographic Characteristics

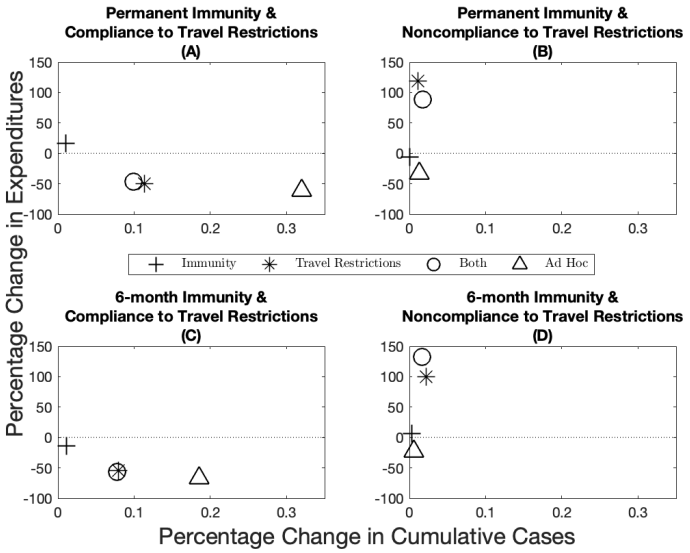

Figure B.13: **Robustness of epidemiological and economic outcomes under different scenarios.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 5%.

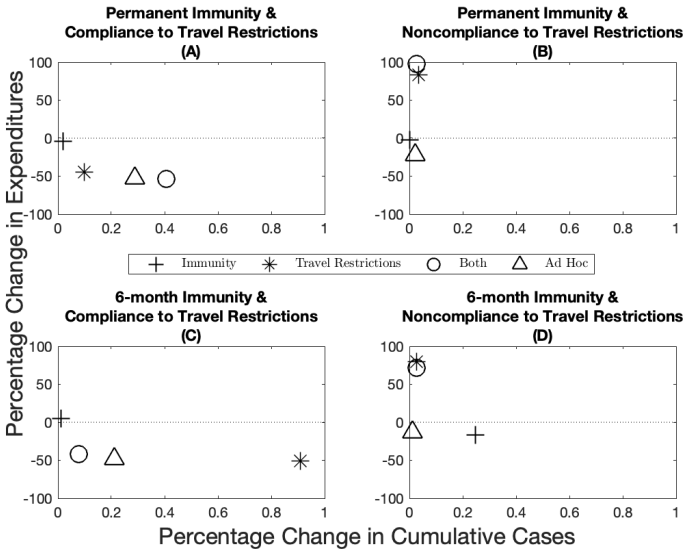

Figure B.14: **Robustness of epidemiological and economic outcomes under different scenarios.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 15%.

B.5.2 Heterogeneous Case-Fatality Ratio

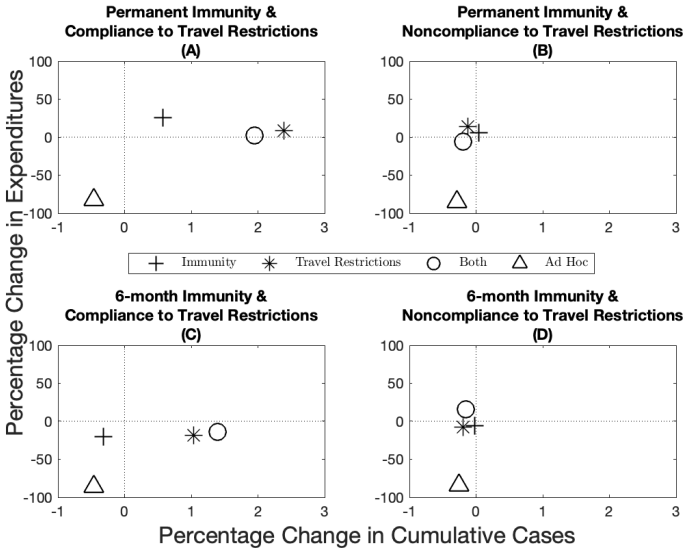

Figure B.15: **Robustness of epidemiological and economic outcomes under different scenarios when the source of heterogeneity is the case-fatality ratio.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 10%.

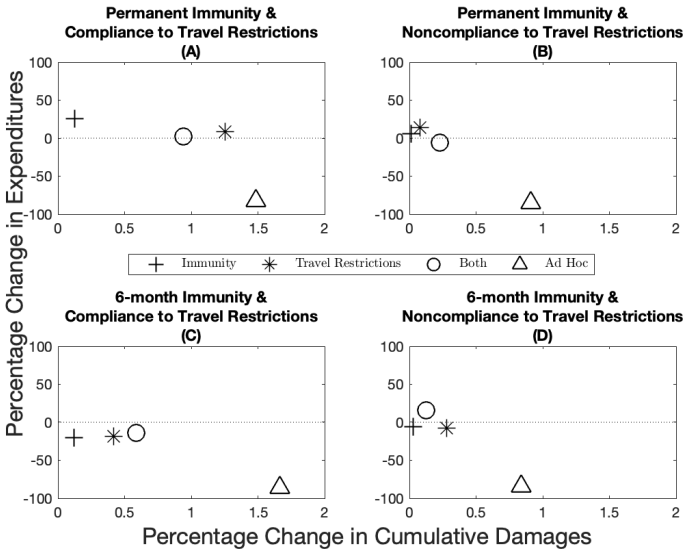

Figure B.16: **Robustness of epidemiological and economic outcomes under different scenarios when the source of heterogeneity is the case-fatality ratio.** Percentage change in expenditures (y-axis) and percentage change in cumulative damages (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 10%. Note that compared to Figure B.15, the use of cumulative damages in this figure gives a more accurate depiction of the situation because cases across jurisdictions are not homogeneous when the case-fatality ratio is different.

B.5.3 Heterogeneous Contact Rate

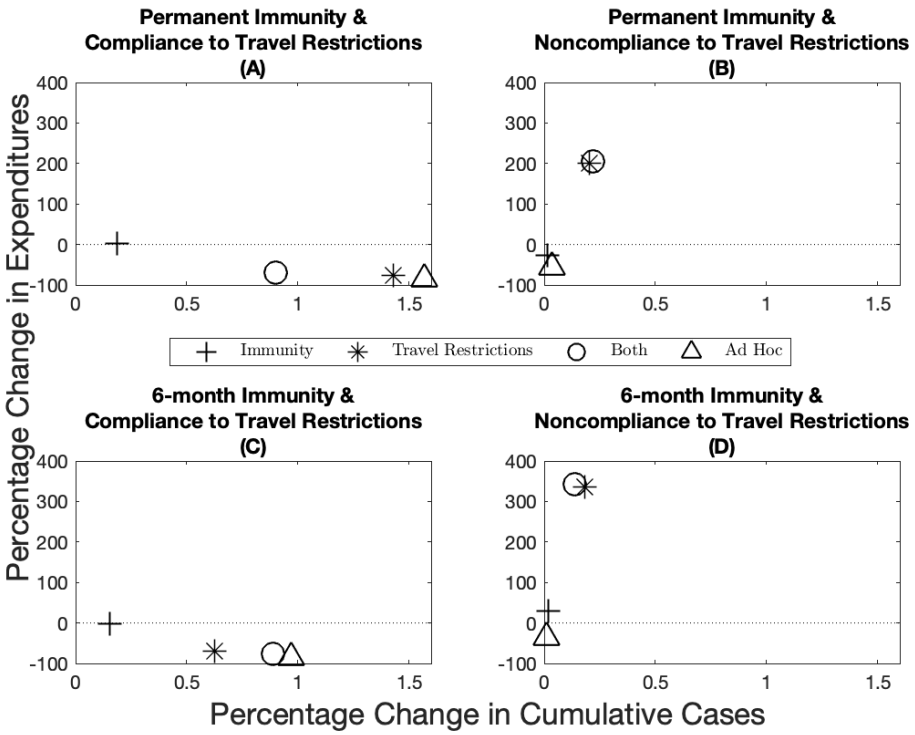

Figure B.17: **Robustness of epidemiological and economic outcomes under different scenarios when the source of heterogeneity is the contact rate.** Percentage change in expenditures (y-axis) and percentage change in cumulative cases (x-axis) from the optimal allocation for different immunity–travel restrictions scenarios and for when vaccine capacity is 10%.

B.6 Figures: Sensitivity Analyses

B.6.1 Workability Cost

As mentioned above, imposing the *ad hoc* rule *ex ante* implicitly means that the central planning agency is essentially assuming that the cost of deviating from the *ad hoc* allocation is infinite. In practice, the workability cost is hard to quantify, because it depends on logistical, political, and cultural factors. It does however seem reasonable to assume, as we did in the paper, that the cost is finite. We investigate the sensitivity of our results by solving for the optimal vaccine allocation over time with levels lower and higher than the base case parameter in the paper. We summarize these results by plotting the variance[35] of the optimal deviation in each time period from the *ad hoc* vaccine allocation (in blue; Figure B.18 panels A, B, C, and D), and the difference in cumulative cases between the optimal and *ad hoc* allocation (in red; Figure B.18 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 B.18 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 B.18 panels E, F, G, and H). If this ratio

[35]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.

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.

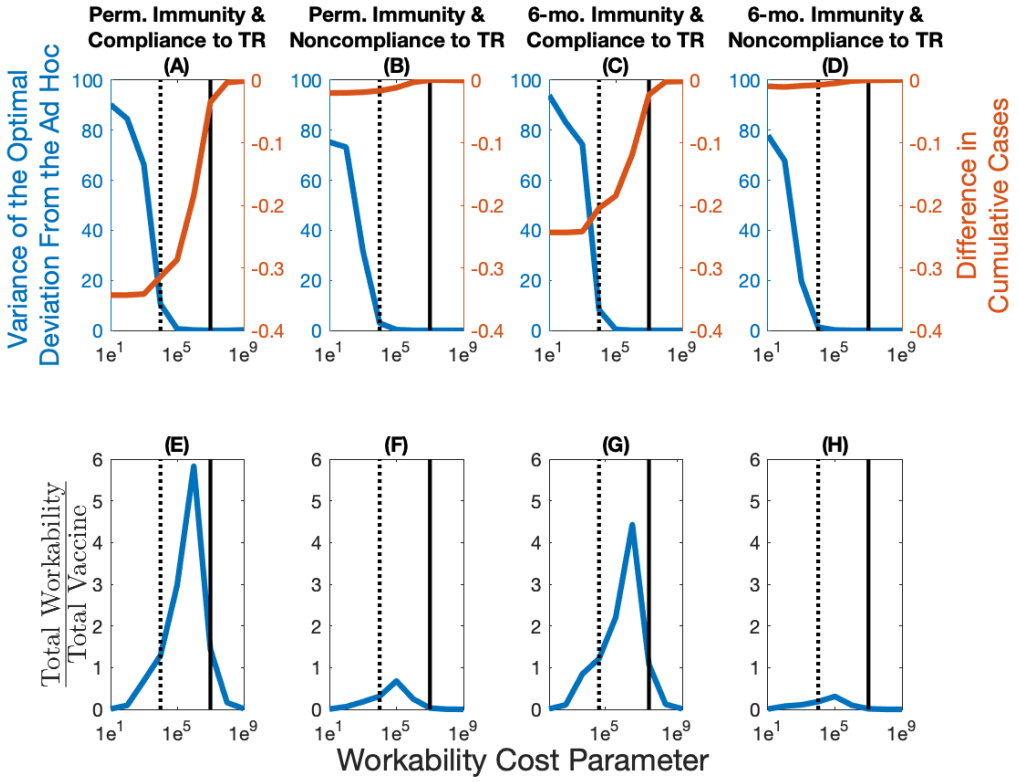

Figure B.18: **Sensitivity of optimal allocations and epidemiological and economic outcomes when varying the workability cost parameter.** The variance of the optimal deviation in percentage (in blue; panels A, B, C, and D) represents an aggregate measure of the optimal deviation from the *ad hoc* allocation. 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 ($1e^7$).

B.6.2 Vaccine Effectiveness

The base case parameter for vaccine effectiveness we utilized in the paper is based on estimates of the influenza vaccine [52]; see Appendix B.1 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 B.19 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 B.19 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 B.19 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.

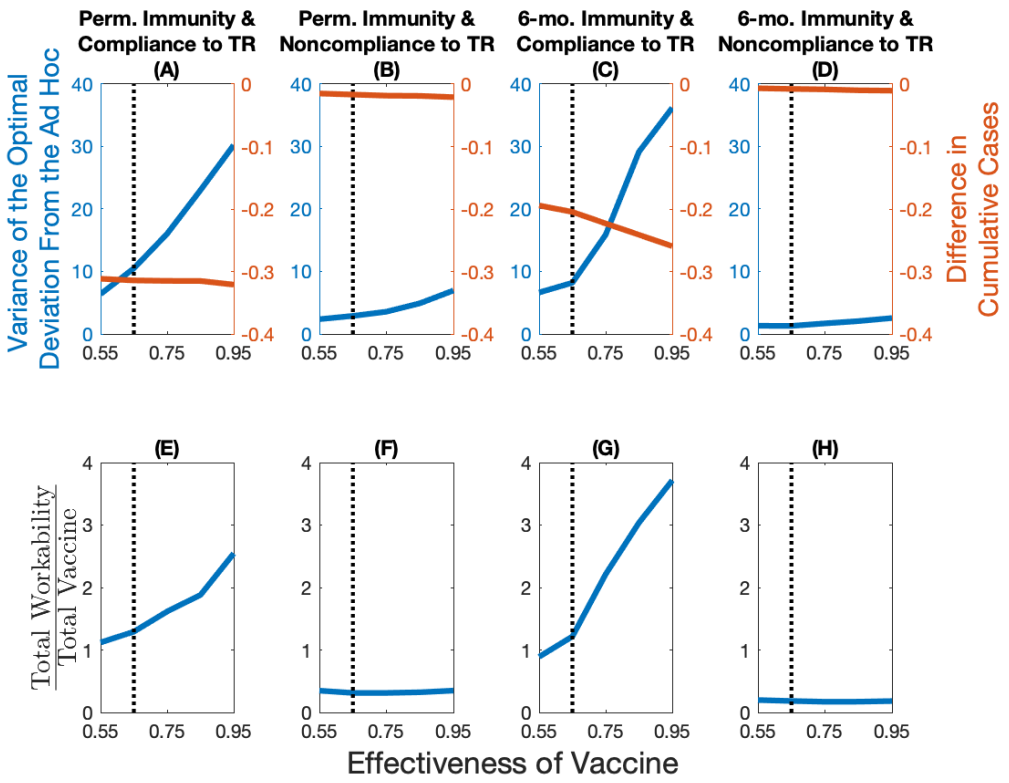

Figure B.19: **Sensitivity of optimal allocations, and of epidemiological and economic outcomes when varying the effectiveness of the vaccine.** The variance of the optimal deviation in percentage (in blue; panels A, B, C, and D) represents an aggregate measure of the optimal deviation from the *ad hoc* 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).

C Chapter 3

C.1 Model & Parameterization

C.1.1 Predator Model

There are two equations of motion describing the dynamics of the prawns over time: (i) the average length of prawns and (ii) the number of prawns. First, the length of prawns—modeled as a somatic growth with the von Bertalanffy growth equation [59]— is given by Equation (11) in the main paper. For the purpose of the Appendix:

$$\dot{L}(t) = k^* \Big(L_\infty - L(t) \Big) \tag{C.1}$$

where $L(t)$ is a state variable denoting the average length of prawns at time t, L_∞ denotes the asymptotic (i.e., maximum) length of prawns, and k^* is given by:

$$k^* = \left(\frac{\max\big\{k, \alpha_N(t)N(t)\big\} + \alpha_U U(t)}{1 + g\Omega(t)} \right)$$

where k is the intrinsic growth parameter, α_i with $i = N, U$ is the attack rate of prawns on snails ($i = N$) and feed ($i = U$), $N(t)$ is the freshwater snail population at time t, $U(t)$ denotes the amount of feeding given to prawns,[36] g is a coefficient parameterized to mimic observed reduction in somatic growth as biomass $\Omega(t)$ increases [2; 61]. Biomass $\Omega(t)$ is defined as the average body size (i.e., the weight) of prawns, $B(t)$, multiplied by the number of prawns, $P(t)$, i.e.,

$$\Omega(t) = B(t)P(t)$$

where body size $B(t)$ relates to length $L(t)$ via the allometric function [62; 63]

$$B(t) = a_p L(t)^{b_p}, \tag{C.2}$$

[36] See [60] for details on feeding of freshwater prawn for aquaculture purposes.

where parameters a_p and b_p are derived by [63].

The number of prawns, $P(t)$, grows according to

$$\dot{P}(t) = -P(t)\Big(\mu_p B(t)^{-d} + \omega\Omega(t)\Big) \tag{C.3}$$

where parameters μ_p, d, and ω are derived from the literature [61; 64; 65]. Equation (C.3) has two components: (i) an exponentially decreasing function of body size (represented by the first term of Equation C.3), and (ii) a linearly increasing function of biomass (represented by the second term of Equation C.3) [2]. The former represents the fact that larger prawns have a lower mortality rate than smaller prawns [65], while the latter represents density–dependent competition for resources and cannibalism at high population densities [61].

C.1.2 Predation Model

Experimental data [66] shows that prawn predation of snails changes with the ratio of prawn body mass to snail body mass. This ratio is given by

$$r(t) = \frac{B(t)}{a_n {L_n}^{b_n}}$$

where the numerator is the prawn body size (given by Equation C.2) and the denominator is an allometric function that relates snail body mass to length (i.e., shell diameter, assumed constant);[37] parameters a_n and b_n are estimated by [2] using data from [66]. Accordingly, the attack rate at time t is defined as [2; 66]:

$$\alpha_N(t) = \alpha_m \log\Big(r(t)\Big)$$

where α_m is a coefficient estimated by [66]. The handling time, i.e., the amount of time it takes the predator to eat the prey, is given by an inverse function of the biomass ratio,

[37] In this paper, we assume a snail length equal to 8mm. Hoover *et al.* [2] divide snail length in three categories: small (4mm), medium (8mm), and large (12mm).

that is,

$$T_h = \Big(T_{hm} r(t)\Big)^{-1}$$

where once again coefficient T_{hm} is estimated by [66]. The per-capita attack rate of prawns on infected snails is given:

$$\psi_W(t) = \left(\frac{\alpha_N(t)\epsilon W(t)^n}{1 + \alpha_N(t)T_{hN}(t)N(t)^n + \alpha_U T_{hU} U(t)^n}\right)$$

where W is the number of infected snails, N is the total number of snails, n is the exponent describing the type of functional response ($n = 2$ for a type III functional response) and ϵ represents additional searching costs for prawns that are not present in laboratory settings where parameters where estimated [2]. Similarly, the per-capita attack rate of prawns on susceptible snails is given by

$$\psi_X(t) = \left(\frac{\alpha_N(t)\epsilon X(t)^n}{1 + \alpha_N(t)T_{hN}(t)N(t)^n + \alpha_U T_{hU} U(t)^n}\right)$$

The predator and predation models are based on field and experimental data [20; 21; 22; 23; 61; 62; 63; 64; 65; 66; 67; 68; 69].

C.1.3 Prey Model

The population of intermediate hosts snails is assumed to grow according to a logistic-type reproduction function [2; 5]. We normalize the carrying capacity of the ecosystem to one such that the intermediate hosts' growth equation is:

$$\dot{X}(t) = fX(t)\Big(1 - N(t)\Big) - \lambda I(t)X(t) - \psi_X(t)P(t) \tag{C.4}$$

where f is the maximal reproduction rate. Note that the nonlinear term of the logistic reproduction function implicitly accounts for density dependent mortality (or fecundity) of susceptible snails.

The maximum reproduction rate of intermediate hosts, f, is set according to estimates of the literature [21], assuming a carrying capacity of 50 individuals per square meter in a water access point of 1,000 square meters [68].

C.1.4 Epidemiological Model

The growth of infected intermediate hosts is given by Equation (14) in the main paper. Note that λ incorporates into a single parameter the rate of reproduction of propagules released in the environment by humans, and the fraction of propagules that reach the environment [7]. The growth of the fraction of infected humans is given by Equation (13) in the main paper.

The basic reproduction number R_0 is defined as the expected number of secondary infection caused by a typical infected individual over its entire infectious period, at a disease-fee equilibrium [16]. Here, to find the R_0 we need to use the next-generation matrix and find its dominant eigenvalue to find the R_0 [16], which in our case turn out to be a function of the shedding rate, the contact rate, and the natural recovery rate. The next-generation matrix is composed of two matrices, denoted F and V, and it is equal to $-FV^{-1}$. 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 [7]. Hence in our model,

$$F = \begin{pmatrix} 0 & \beta \\ 0 & 0 \end{pmatrix} \text{ and } V = \begin{pmatrix} -\gamma & 0 \\ \lambda & -\delta \end{pmatrix}$$

where the two rows of F and V refer to the $I(t)$ and $W(t)$ equations, respectively. Note that matrices are derived under the assumption of introducing a single infected snail in an otherwise susceptible population. Following the methodology of Diekmann *et al.* [16], the basic reproduction ratio is given by,

$$R_0 = \frac{\beta\lambda}{\delta\gamma}.$$

Following Castonguay *et al.* [57], the contact rate β is set such that there is 1 infection per 200 water contacts (which is consistent with findings of the literature [19; 68]) and the shedding rate λ is set such that there is approximately 1 intermediate host infection per 555 sheds. Parameter values β and λ were calibrated to match $R_0 = 3.5$; note that according to Sokolow *et al.* [68], the expected R_0 for schistosomiasis ranges from 1 to 7. Parameter γ, the natural recovery rate in humans, is calibrated using the life expectancy of the disease in hosts (3.3 years) [20], and assuming that humans, once infected, carry 70 worms [68]. The natural mortality rate of intermediate hosts, δ, is based on their life expectancy (2 months) [21; 22; 23].

C.1.5 Economic Model

Damages due to infection occur on the human population; let $c_I(I)$ be the damage function. Because the health benefits greatly outweigh the aquaculture profits, if we were to include the integral of costs, i.e. "$\int_0^\infty c_I(I)\ \mathrm{d}t$", in the objective function, then the solution would be to choose the rotation length to be as small as possible to minimize infection costs. We would need to incorporate some form of scrap value associated with infected humans, but programming complications would arise because some state variables (i.e. L and P) would start anew each rotation, while other state variables (i.e. I, W, and X) have a memory process from one rotation to the other. As a reason, we opted to minimize the avoided infection costs, which approximates very well what we were looking for: minimizing the health impact. Hence, we assume that the damage function is:

$$c_I(I) = c_I\Big[I(0) - e^{-rT} I(T)\Big]$$

where c_I is the cost parameter associated with infected humans. This parameter represents disability and reduced intellectual function [17] causing lower school participation for children [70] and lower worker productivity for adults [71; 72; 73; 74].

The other important economic component of the model is the feeding function, which we model as a quadratic function in the amount of feed utilized:

$$c_u\big(U(t)\big) = c_u U(t)^2$$

where c_u is a cost parameter associated with the feeding cost function.

Benefits and costs are discounted using a discount rate of $r = 0.07$ or 7%, mimicking relatively high discount rates in sub-Saharan Africa [2]. The damage parameter c_I is calibrated using estimates of Lo *et al.* [5]. They find that, without intervention, infection rates of 38% (our steady-state initial conditions) in a community of 5,000 people would yield losses of 550 disability-adjusted life years (DALYs). Following Castonguay *et al.* [57], the value of a life year was set to be approximately the median value of the GDP per capita of an African country (approximately $3,000 USD). The feeding cost function is based on estimates of feed costs from the aquaculture literature [69] and estimates of transportation cost estimates from the schistosomiasis literature [4; 5; 11; 12; 13; 14; 15].

C.1.6 Summary of Parameter Levels

Parameters	Level	Interpretation
L_∞	213.63	Asymptotic length of prawns (in mm).[38]
k	3.796	Maximum growth rate of prawns (in mm/year).[39]
g	3.50×10^{-6}	Density-dependent growth reduction.[40]
a_P	0.0733	Allometric parameters for prawns length-weight relationship.[41]
b_P	3.5502	Allometric parameters for prawns length-weight relationship.[42]
μ_P	2.21	Natural prawn mortality rate.[43]
d	-0.382	Size-dependent mortality scaling coefficient.[44]
ω	5.50×10^{-9}	Density dependent mortality factor.[45]
a_n	0.1872	Allometric parameters for snails length-weight relationship.[46]
b_n	2.5368	Allometric parameters for snail length-weight relationship.[47]
α_m	0.9050	Coefficient for relationship between biomass ratio and attack rate.[48]
T_{hm}	0.38561	Coefficient for relationship between biomass ratio and handling time.[49]
n	2	Exponent of Holling's type III functional response.[50]
ϵ	0.1	Prawn predation attack rate penalty.[51]
f	1.168	Maximum reproduction rate of intermediate hosts.[52]
δ	1.20×10^{-4}	Natural death rate of the disease in the environment.[53]
β	5.20×10^{-3}	Contact rate of humans.[54]
λ	3.50×10^{-4}	Shedding rate of humans.[55]
γ	4.30×10^{-3}	Natural recovery rate of humans.[56]
p	12	Price of prawns (per kg).[57]
c_P	0.1	Cost of juvenile prawns.[58]
c_U	200	Cost parameter associated with prawn feeding.[59]
r	0.07	Discount rate.[60]

Table C.1: Parameter levels used in the numerical simulation.

C.2 Optimization

C.2.1 Boundary Conditions

The initial and terminal conditions on the length and the number of prawns are respectively given by:

$$L(0) \text{ is given, and } L(T) \text{ is free;} \tag{C.5a}$$

$$P(0) \text{ is given, and } P(T) \text{ is free;} \tag{C.5b}$$

$$W(0) \text{ is given, and } W(T) \text{ is free;} \tag{C.5c}$$

$$X(0) \text{ is given, and } X(T) \text{ is free;} \tag{C.5d}$$

$$I(0) \text{ is given, and } I(T) \text{ is free;} \tag{C.5e}$$

where it is assumed that $L(0)$ and $P(0)$ are the same values as the ones used in Hoover *et al.* [2], and $W(0)$, $X(0)$ and $I(0)$ are at their pre-treatment equilibria; $L(T)$, $P(T)$, $W(T)$, $X(T)$, and $I(T)$ are optimally chosen, conditional on $T > 0$.

[38] Based on estimates from [64].
[39] Based on estimates from [62].
[40] Based on estimates from [2; 61].
[41] Based on estimates from [63].
[42] *Ibid.*
[43] Based on estimates from [64].
[44] Based on estimates from [65].
[45] Based on estimates from [2; 61].
[46] Based on experimental data of [66] and estimated by [2].
[47] *Ibid.*
[48] *Ibid.*
[49] *Ibid.*
[50] Based on [67].
[51] Meant to represent additional searching cost for prey in wildlife rather than laboratory conditions where the data was collected [66]; based on [2].
[52] Based on an instantaneous intrinsic fertility rate of snails of 0.16 per day per square meter [68], and assuming a carrying capacity of 50,000 individuals in a water access point of 1,000 square meters [2].
[53] Based on the life expectancy of infected intermediate hosts (2 months, 21; 22; 23), assuming a carrying capacity of 50,000 individuals in a water access point of 1,000 square meters [2].
[54] Calibrated to match a R_0 of 3.5 [68].
[55] *Ibid.*
[56] Based on the life expectancy of the disease in hosts, i.e., 3.3 years [20], assuming humans carry 70 worms once infected [68].
[57] Based on estimates from [69].
[58] *Ibid.*
[59] *Ibid.*
[60] Meant to mimic relatively high discount rates in sub-Saharan Africa [2].

C.2.2 Nonnegativity and Upper-Bound Constraints

State variables $L(t)$ and $P(t)$ are also subject to constraints that bound them from above and below. Formally, the constraints are given by:

$$0 \leq L(t) \leq L_{\infty}; \tag{C.6a}$$

$$0 \leq P(t) \leq P(0) \tag{C.6b}$$

$$0 \leq W(t) \leq 1 \tag{C.6c}$$

$$0 \leq X(t) \leq 1 \tag{C.6d}$$

$$0 \leq I(t) \leq 1 \tag{C.6e}$$

where the upper bound in Equation (C.6a) follows since the prawns cannot grow more than their asymptotic length, the upper bound in Equation and (C.6b) follows since the number of prawns in the enclosure exhibits a negative density-dependent competition for resources and cannibalism at high population densities, where the upper bounds in equations (C.6c) and (C.6d) follow since the carrying capacity of the ecosystem was normalize to one, and finally, where the upper bound on (C.6e) follows since $I(t)$ represents the fraction of humans that are infected. The control variable representing the quantity of supplemental feed provided to prawns is only bounded by a nonnegativity constraint:

$$0 \leq U(t). \tag{C.7}$$

C.2.3 Objective Functions

When the farmer only cares about the benefits and costs of one single rotation and it does not feed prawns, the private objective function is given in the main paper by Equation (15). In practice, it is likely that the decision-maker, whether its objective is health or profit motivated, will care about several rotations given the relatively short time of the rotation. When the farmer cares about the benefits and costs of an infinite number of rotations and it has the opportunity to supplement prawns' diet with feed, the private

objective is:

$$\max_{T,U(t)} \left(\frac{1}{e^{rT}-1}\right)\left(p\Omega(T) - c_P P(0) - \int_0^T e^{-rt} c_U\big(U(t)\big)\ \mathrm{d}t\right) - c_P P(0) - \int_0^T e^{-rt} c_U\big(U(t)\big)\ \mathrm{d}t. \tag{C.8}$$

The health-maximizing objective when the planner cares about the benefits and costs of one single rotation is given by Equation (16) in the main paper. As with the private-maximzing objective, the decision-maker is likely to care about the benefits and costs of an infinite number of rotations. In such a case, the objective function is:

$$\max_{T,U(t)} \left(\frac{1}{e^{rT}-1}\right)\left(c_I\Big[I(0) - I(T)\Big] - c_P P(0) - \int_0^T e^{-rt} c_U\big(U(t)\big)\ \mathrm{d}t\right) - c_P P(0) - \int_0^T e^{-rt} c_U\big(U(t)\big)\ \mathrm{d}t + c_I I(0). \tag{C.9}$$

The above objective functions are subject to predators' dynamics (equations 11 and C.3), the preys' dynamics (Equation C.4), the dynamics of disease transmission (equations 13 and 14), along with free endpoint conditions (equations C.5), non-negativity and physical constraints on the predators and preys (equations C.6), and the non-negativity constraint on feed (Equation C.7).

C.2.4 Numerical Methods

In this paper, we use pseudospectral collocation to approximate the continuous time problem with a constrained nonlinear programming problem (see [31; 32; 57; 58; 75] for other applications of this technique). The dynamic control to our problem—i.e. the quantity of feed applied—are approximated by a polynomial of degree n (determined by the number of collocation points) over the endogenously determined time horizon [34]. The residual error of the constraints (e.g. the nonnegativity constraints) is minimized by the algorithm at the n collocation points, where we chose n to have a reasonable speed of convergence to a solution and a low numerical error. Here, we chose 60 collocations

points. In problem like the one studied in this paper, 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) on state and control variables can be directly incorporated in the optimal control problem [35]. Pseudospectral collocation enables us to potentially obtain a solution that may lay on the boundary of the control set. Because we are designing policies that have dynamic controls (e.g. the limited feeding season), this is necessarily going to be the case because feed will be ban for some part of the time horizon. Another advantage of this method is the ability to deal with large-scale dynamical systems, such as the one presented here with five state variables, one control variable, and one endogenously determined time horizon. The solution was found using TOMLAB (v. 8.4) [36; 37] and the accompanying PROPT toolbox [38]. We used the general-purpose nonlinear optimization packages KNITRO, SNOPT and NPSOL to solve the different approximated nonlinear programming problems.

C.3 Figures: Varying Feed Conversion Efficiency

C.3.1 Aquaculture Profits by Policy

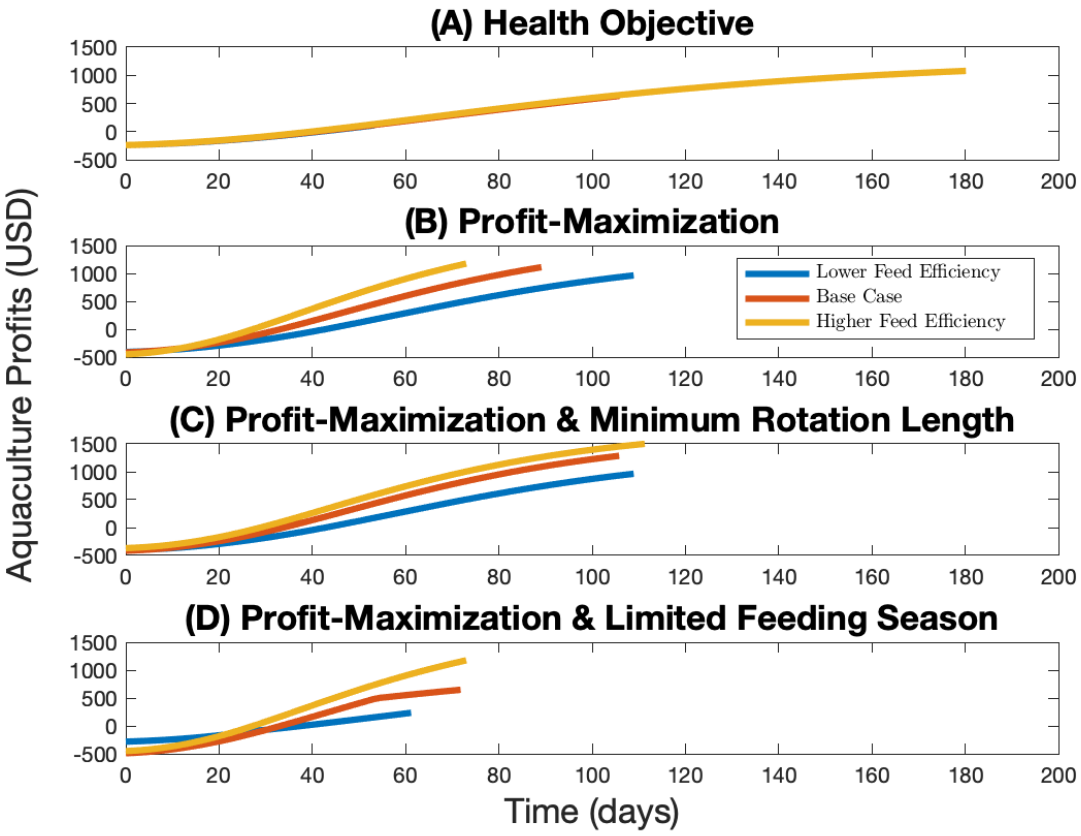

Figure C.1: **Dynamics of the aquaculture profit functions under different policies and for different feed conversion efficiencies.** The figure depicts the change in harvesting profits when health outcomes are maximized (Panel A), when profit is maximized and there is no policy (Panel B), when profit is maximized with a minimum rotation length policy (Panel C), and when profit is maximized with a limited feeding season policy (Panel D) for a lower feed conversion efficiency (i.e. 50% lower, in blue), the base case feed conversion efficiency (in red), and for a higher feed conversion efficiency (i.e. 50% higher, in yellow).

C.3.2 Infected Snails by Policy

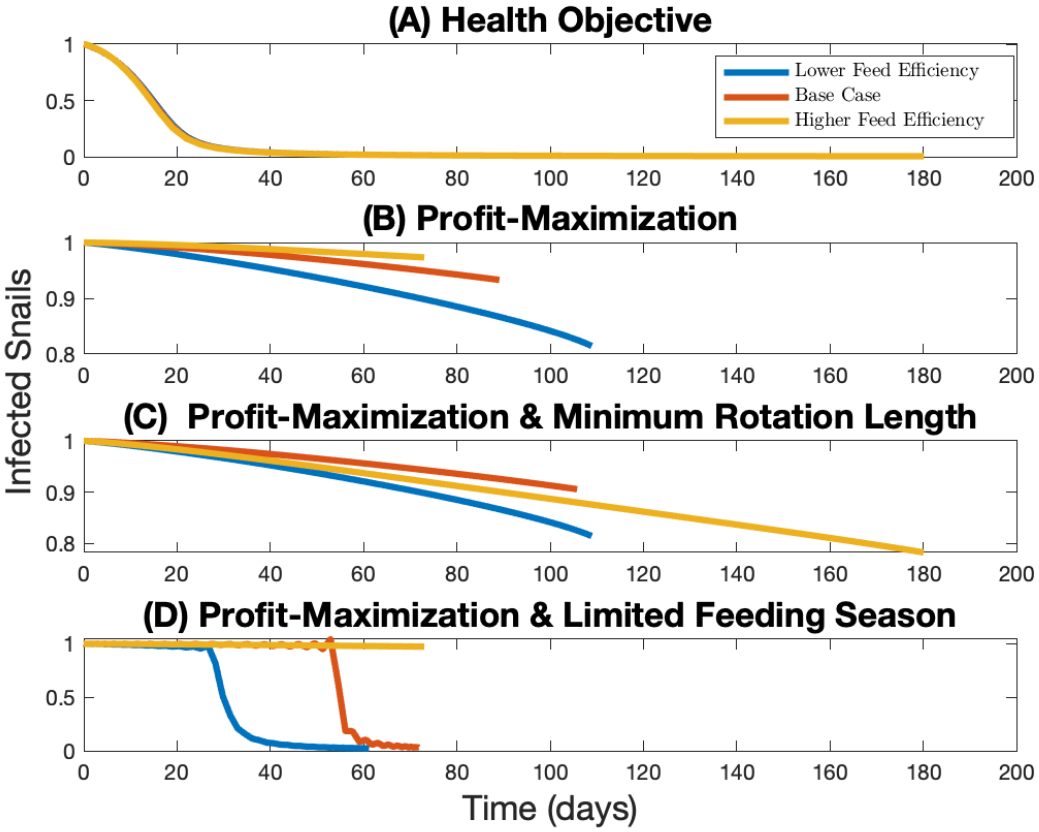

Figure C.2: **Dynamics of the infected snail population under different policies and for different feed conversion efficiencies.** The figure depicts the change over time in the number of infected snails when health outcomes are maximized (Panel A), when profit is maximized and there is no policy (Panel B), when profit is maximized with a minimum rotation length policy (Panel C), and when profit is maximized with a limited feeding season policy (Panel D) for a lower feed conversion efficiency (i.e. 50% lower, in blue), the base case feed conversion efficiency (in red), and for a higher feed conversion efficiency (i.e. 50% higher, in yellow).

C.3.3 Quantity of Feed by Policy

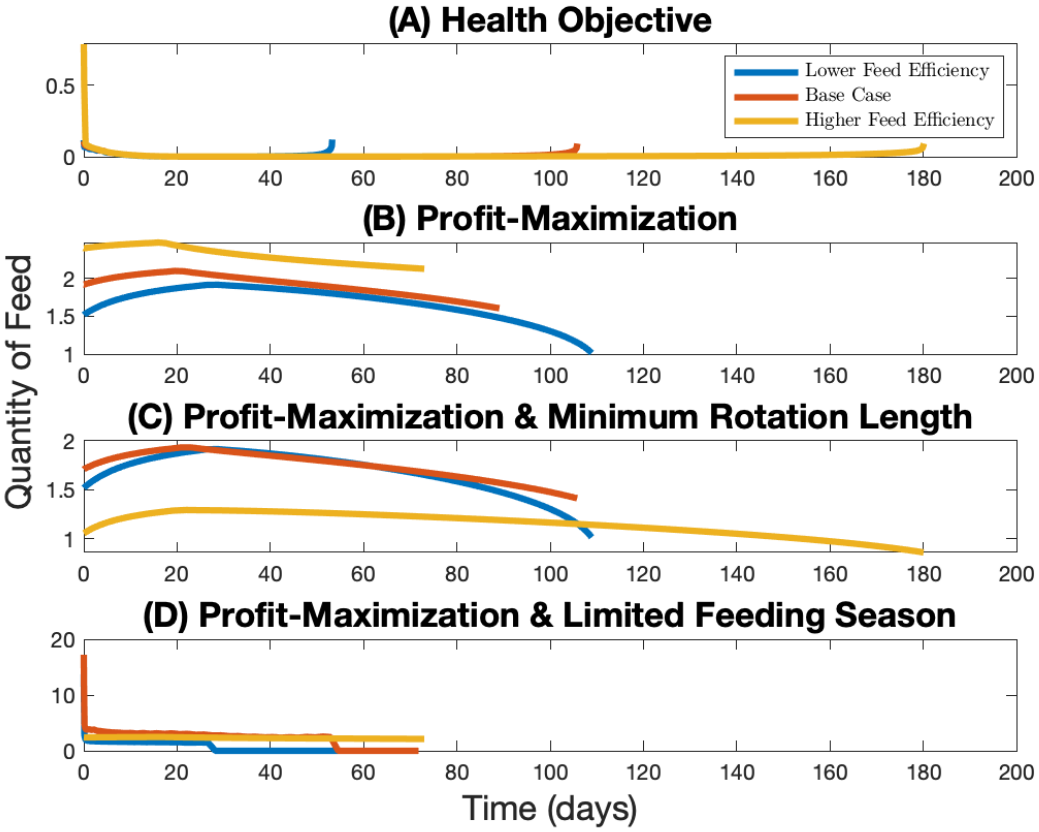

Figure C.3: **Dynamics of the feeding paths under different policies and for different feed conversion efficiencies.** The figure depicts the change over time in feed when health outcomes are maximized (Panel A), when profit is maximized and there is no policy (Panel B), when profit is maximized with a minimum rotation length policy (Panel C), and when profit is maximized with a limited feeding season policy (Panel D) for a lower feed conversion efficiency (i.e. 50% lower, in blue), the base case feed conversion efficiency (in red), and for a higher feed conversion efficiency (i.e. 50% higher, in yellow).

C.4 Figures: Varying The Holling Type III Functional Response Exponent

C.4.1 Aquaculture Profits: Health- and Profit-Maximizing Outcomes

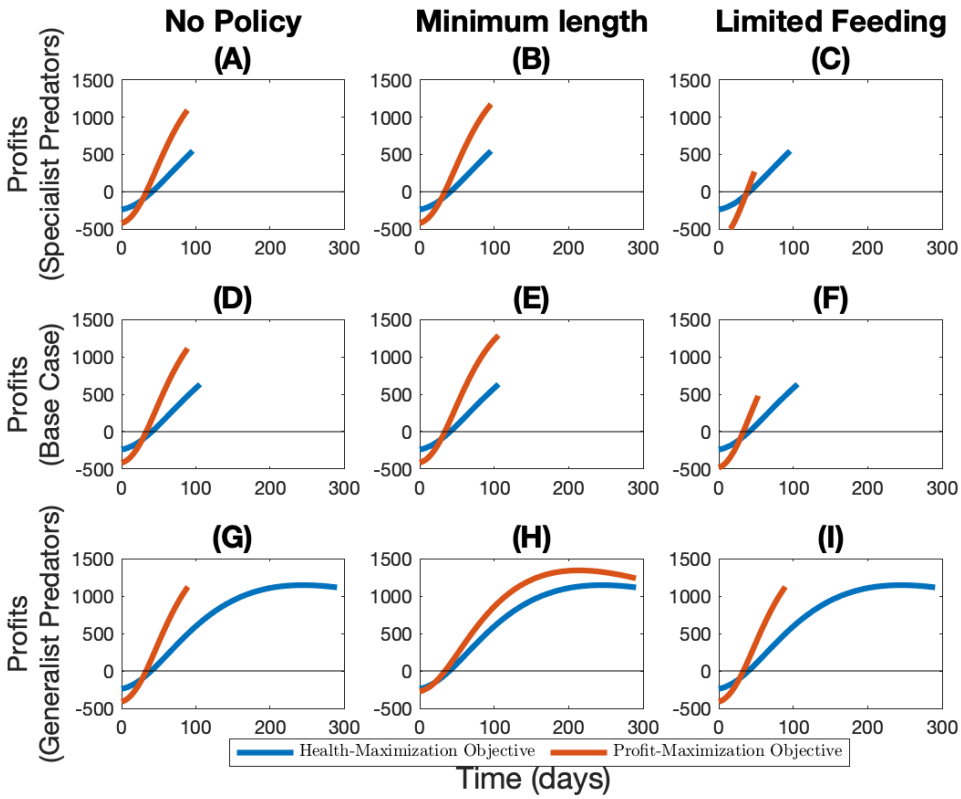

Figure C.4: **Dynamics of the aquaculture profit functions under different policy and for different Holling type III functional response exponent.** The figure depicts the change in harvesting profits when there is no policy (panels A, D, and G), with a minimum rotation length policy (panels B, E, and H), and with a limited feeding season policy (panels C, F, and I) for the health- (in blue) and profit-maximizing (in red) objectives. Panels A, B, and C represent the case where prawns have a lower Holling type III functional response exponent than in the base case (i.e. $n = 1.75$); panels D, E, and F represent the base case Holling type III functional response exponent (i.e. $n = 2$); panels G, H, and I represent a case where prawns have a higher Holling type III functional response exponent than in the base case (i.e. $n = 2.25$).

C.4.2 Infected Snails: Health- and Profit-Maximizing Outcomes

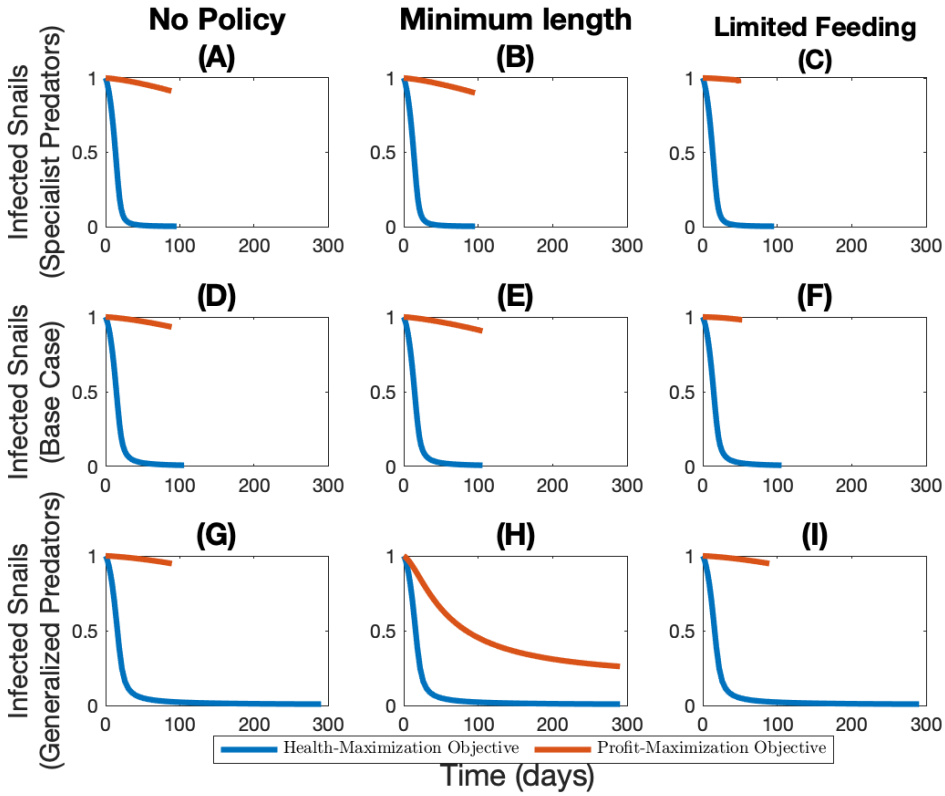

Figure C.5: **Dynamics of the infected snail population under different policy and for different Holling type III functional response exponent.** The figure depicts the change over time in the number of infected snails as a proportion of its steady-state value when there is no policy (panels A, D, and G), with a minimum rotation length policy (panels B, E, and H), and with a limited feeding season policy (panels C, F, and I) for the health- (in blue) and profit-maximizing (in red) objectives. Panels A, B, and C represent the case where prawns have a lower Holling type III functional response exponent than in the base case (i.e. $n = 1.75$); panels D, E, and F represent the base case Holling type III functional response exponent (i.e. $n = 2$); panels G, H, and I represent a case where prawns have a higher Holling type III functional response exponent than in the base case (i.e. $n = 2.25$).

C.4.3 Feeding Paths: Health- and Profit-Maximizing Outcomes

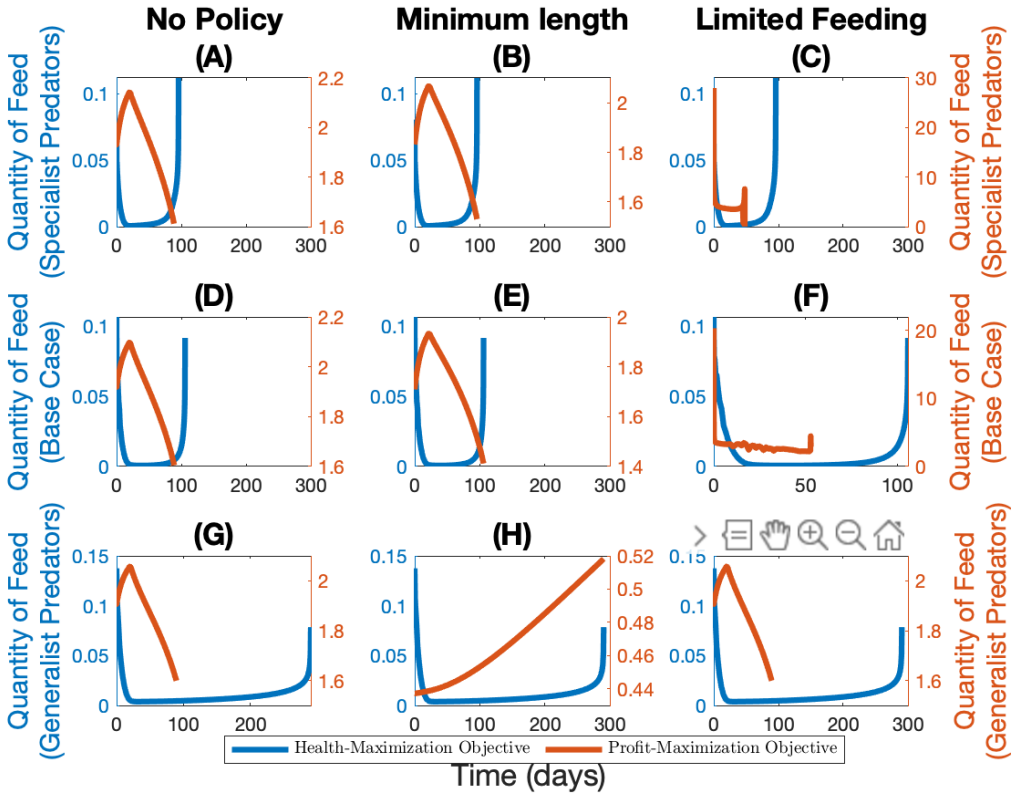

Figure C.6: **Dynamics of the feeding paths under different policy and for different Holling type III functional response exponent.** The figure depicts the change over time in feed when there is no policy (panels A, D, and G), with a minimum rotation length policy (panels B, E, and H), and with a limited feeding season policy (panels C, F, and I) for the health- (in blue) and profit-maximizing (in red) objectives. Panels A, B, and C represent the case where prawns have a lower Holling type III functional response exponent than in the base case (i.e. $n = 1.75$); panels D, E, and F represent the base case Holling type III functional response exponent (i.e. $n = 2$); panels G, H, and I represent a case where prawns have a higher Holling type III functional response exponent than in the base case (i.e. $n = 2.25$).

C.4.4 Aquaculture Profits by Policy

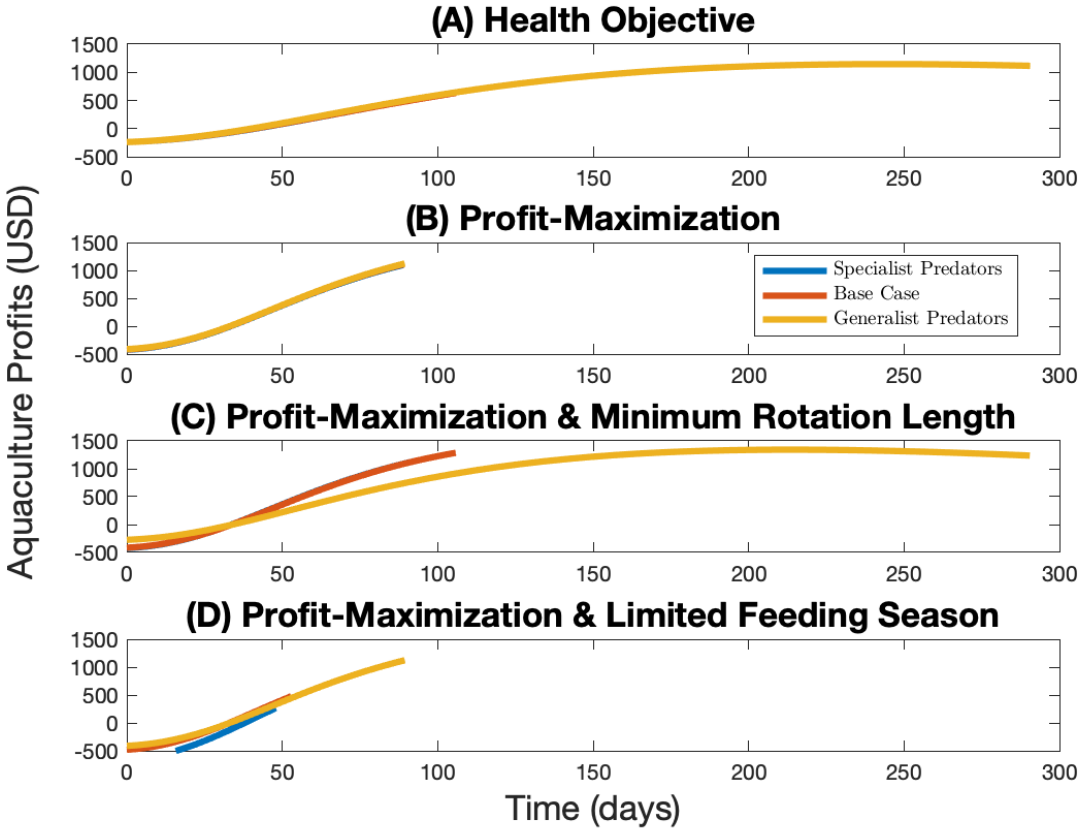

Figure C.7: **Dynamics of the aquaculture profit functions under different policies and for different Holling type III functional response exponent.** The figure depicts the change in harvesting profits when health outcomes are maximized (Panel A), when profit is maximized and there is no policy (Panel B), when profit is maximized with a minimum rotation length policy (Panel C), and when profit is maximized with a limited feeding season policy (Panel D) for a lower feed conversion efficiency (i.e. 50% lower, in blue), the base case feed conversion efficiency (in red), and for a higher feed conversion efficiency (i.e. 50% higher, in yellow).

C.4.5 Infected Snails by Policy

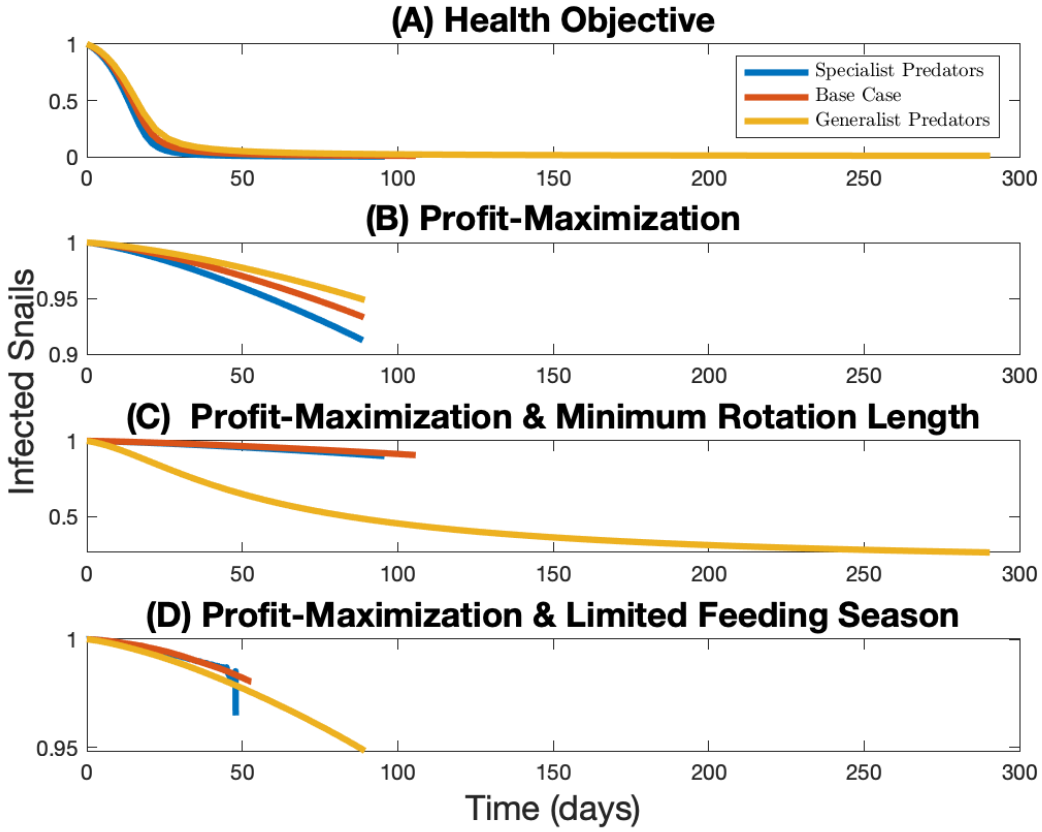

Figure C.8: **Dynamics of the infected snail population under different policies and for different Holling type III functional response exponent.** The figure depicts the change over time in the number of infected snails when health outcomes are maximized (Panel A), when profit is maximized and there is no policy (Panel B), when profit is maximized with a minimum rotation length policy (Panel C), and when profit is maximized with a limited feeding season policy (Panel D) for a lower feed conversion efficiency (i.e. 50% lower, in blue), the base case feed conversion efficiency (in red), and for a higher feed conversion efficiency (i.e. 50% higher, in yellow).

C.4.6 Quantity of Feed by Policy

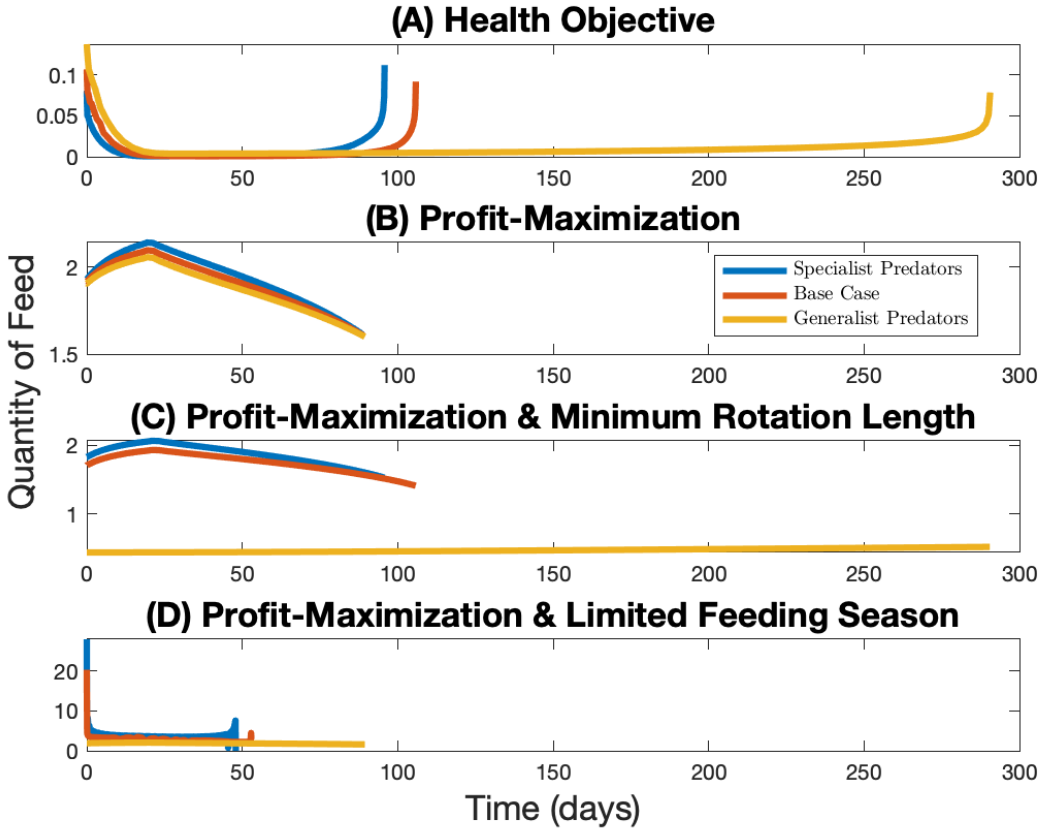

Figure C.9: **Dynamics of the feeding paths under different policies and for different Holling type III functional response exponent.** The figure depicts the change over time in feed when health outcomes are maximized (Panel A), when profit is maximized and there is no policy (Panel B), when profit is maximized with a minimum rotation length policy (Panel C), and when profit is maximized with a limited feeding season policy (Panel D) for a lower feed conversion efficiency (i.e. 50% lower, in blue), the base case feed conversion efficiency (in red), and for a higher feed conversion efficiency (i.e. 50% higher, in yellow).

References

[1] Klepac P, Laxminarayan R, Grenfell BT. Synthesizing epidemiological and economic optima for control of immunizing infections. Proceedings of the National Academy of Sciences. 2011;108(34):14366–14370.

[2] Hoover CM, Sokolow SH, Kemp J, Sanchirico JN, Lund AJ, Jones IJ, et al. Modelled effects of prawn aquaculture on poverty alleviation and schistosomiasis control. Nature Sustainability. 2019;2(7):611.

[3] Sokolow SH, Huttinger E, Jouanard N, Hsieh MH, Lafferty KD, Kuris AM, et al. Reduced transmission of human schistosomiasis after restoration of a native river prawn that preys on the snail intermediate host. Proceedings of the National Academy of Sciences. 2015;112(31):9650–9655.

[4] Lo NC, Bogoch II, Blackburn BG, Raso G, N'Goran EK, Coulibaly JT, et al. Comparison of community-wide, integrated mass drug administration strategies for schistosomiasis and soil-transmitted helminthiasis: a cost-effectiveness modelling study. The Lancet Global Health. 2015;3(10):e629–e638.

[5] Lo NC, Gurarie D, Yoon N, Coulibaly JT, Bendavid E, Andrews JR, et al. Impact and cost-effectiveness of snail control to achieve disease control targets for schistosomiasis. Proceedings of the National Academy of Sciences. 2018;115(4):E584–E591.

[6] Nguyen C, Carlson JM. Optimizing real-time vaccine allocation in a stochastic SIR model. PloS one. 2016;11(4):e0152950.

[7] Garchitorena A, Sokolow S, Roche B, Ngonghala C, Jocque M, Lund A, et al. Disease ecology, health and the environment: a framework to account for ecological and socio-economic drivers in the control of neglected tropical diseases. Philosophical Transactions of the Royal Society B: Biological Sciences. 2017;372(1722):20160128.

[8] Yang GJ, Li W, Sun LP, Wu F, Yang K, Huang YX, et al. Molluscicidal efficacies of different formulations of niclosamide: result of meta-analysis of Chinese literature. Parasites & Vectors. 2010;3(1):84.

[9] Ross A, Bartley P, Sleigh A, Olds G, Li Y, Williams G, et al. Schistosomiasis. New England Journal of Medicine. 2002;346:1212–1220.

[10] Hotez PJ, Fenwick A, Kjetland EF. Africa's 32 cents solution for HIV/AIDS. PLoS Neglected Tropical Diseases. 2009;3(5):e430.

[11] World Health Organization. Preventive chemotherapy in human helminthiasis. Coordinated use of anthelminthic drugs in control interventions: a manual for health professionals and programme managers. World Health Organization; 2006.

[12] Barnish G. Evaluation of Chemotherapy in the Control of Schistosoma Mansoni in Marquis Valley, Saint Lucia. The American Journal of Tropical Medicine and Hygiene. 1982;31(1):111–115.

[13] King CH, Olbrych SK, Soon M, Singer ME, Carter J, Colley DG. Utility of repeated praziquantel dosing in the treatment of schistosomiasis in high-risk communities in Africa: a systematic review. PLoS Neglected Tropical Diseases. 2011;5(9):e1321.

[14] Guyatt H. The cost of delivering and sustaining a control programme for schistosomiasis and soil-transmitted helminthiasis. Acta Tropica. 2003;86(2-3):267–274.

[15] Barnish G, Jordan P, Bartholomew RK, Grist E. Routine focal molluscicide after chemotherapy to control Schistosoma mansoni in Cul de Sac valley, Saint Lucia. Transactions of the Royal Society of Tropical Medicine and Hygiene. 1982;76(5):602–609.

[16] Diekmann O, Heesterbeek JAP, Metz JA. On the definition and the computation of the basic reproduction ratio R0 in models for infectious diseases in heterogeneous populations. Journal of Mathematical Biology. 1990;28(4):365–382.

[17] Colley DG, Bustinduy AL, Secor WE, King CH. Human schistosomiasis. The Lancet. 2014;383(9936):2253–2264.

[18] Inobaya MT, Olveda RM, Chau TN, Olveda DU, Ross AG. Prevention and control of schistosomiasis: a current perspective. Research and Reports in Tropical Medicine. 2014;2014(5):65.

[19] Woolhouse M, Mutapi F, Ndhlovu P, Chandiwana S, Hagan P. Exposure, infection and immune responses to Schistosoma haematobium in young children. Parasitology. 2000;120(1):37–44.

[20] Goddard M, Jordan P. On the longevity of Schistosoma mansoni in man on St. Lucia, West Indies. Transactions of the Royal Society of Tropical Medicine and Hygiene. 1980;74(2):185–191.

[21] Woolhouse M, Chandiwana S. Population biology of the freshwater snail Bulinus globosus in the Zimbabwe highveld. Journal of Applied Ecology. 1990:41–59.

[22] Kariuki HC, Ivy JA, Muchiri EM, Sutherland LJ, King CH. Divergent effects of Schistosoma haematobium exposure on intermediate-host snail species Bulinus nasutus and Bulinus globosus from coastal Kenya. The American Journal of Tropical Medicine and Hygiene. 2017;96(4):850–855.

[23] Anderson R, May R. Prevalence of schistosome infections within molluscan populations: observed patterns and theoretical predictions. Parasitology. 1979;79(1):63–94.

[24] Central Intelligence Agency. The World Factbook. Washington, DC: Central Intelligence Agency, 2018. `https://www.cia.gov/library/publications/the-world-factbook/index.html` (accessed November 10, 2018); 2018.

[25] World Health Organization. Accelerating work to overcome the global impact of neglected tropical diseases: a roadmap for implementation: executive summary. Geneva: World Health Organization; 2012.

[26] World Health Organization. Handbook for integrated vector management. Geneva: World Health Organization; 2012.

[27] Dyson L, Stolk WA, Farrell SH, Hollingsworth TD. Measuring and modelling the effects of systematic non-adherence to mass drug administration. Epidemics. 2017;18:56–66.

[28] Utzinger J, Zhou XN, Chen MG, Bergquist R. Conquering schistosomiasis in China: the long march. Acta Tropica. 2005;96(2-3):69–96.

[29] Guo JG, Chun-Li C, Guang-Han H, Han L, Dong L, Rong Z, et al. The role of 'passive chemotherapy'plus health education for schistosomiasis control in China during maintenance and consolidation phase. Acta Tropica. 2005;96(2-3):177–183.

[30] Gurarie D, King CH, Yoon N, Li E. Refined stratified-worm-burden models that incorporate specific biological features of human and snail hosts provide better estimates of Schistosoma diagnosis, transmission, and control. Parasites & Vectors. 2016;9(1):428.

[31] Sanchirico JN, Springborn M. How to get there from here: ecological and economic dynamics of ecosystem service provision. Environmental and Resource Economics. 2011;48(2):243–267.

[32] Kling DM, Sanchirico JN, Wilen JE. Bioeconomics of managed relocation. Journal of the Association of Environmental and Resource Economists. 2016;3(4):1023–1059.

[33] Fuller KB, Sanchirico JN, Alston JM. The spatial-dynamic benefits from cooperative disease control in a perennial crop. Journal of Agricultural and Resource Economics. 2017;42(2):127–145.

[34] Garg D, Patterson M, Hager WW, Rao AV, Benson DA, Huntington GT. A unified framework for the numerical solution of optimal control problems using pseudospectral methods. Automatica. 2010;46(11):1843–1851.

[35] Judd KL. Numerical methods in economics. MIT press; 1998. Cambridge, Massachusetts.

[36] Holmström K. Practical optimization with the tomlab environment in matlab. In: Proceedings of the 42nd SIMS Conference. Citeseer; 2001 Sep 15. p. 89–108. Porsgrunn, Norway: Telemark University College.

[37] Holmström K, Göran AO, Edvall MM. Users guide for TOMLAB/SNOPT. Mälardalen University, Department of Mathematics and Physics, Västerås, Sweden. 2008.

[38] Rutquist PE, Edvall MM. Propt-Matlab optimal control software. Tomlab Optimization Inc. 2010;260(1).

[39] Sheldon TA. Discounting in health care decision-making: time for a change? Journal of Public Health. 1992;14(3):250–256.

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

[41] Diekmann O, Heesterbeek J, Roberts MG. The construction of next-generation matrices for compartmental epidemic models. Journal of the Royal Society Interface. 2010;7(47):873–885.

[42] Li Q, Guan X, Wu P, Wang X, Zhou L, Tong Y, et al. Early transmission dynamics in Wuhan, China, of novel coronavirus–infected pneumonia. New England Journal of Medicine. 2020.

[43] Tian H, Liu Y, Li Y, Wu CH, Chen B, Kraemer MU, et al. An investigation of transmission control measures during the first 50 days of the COVID-19 epidemic in China. Science. 2020;368(6491):638–642.

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

[45] Abdollahi E, Champredon D, Langley JM, Galvani AP, Moghadas SM. Temporal estimates of case-fatality rate for COVID-19 outbreaks in Canada and the United States. CMAJ. 2020.

[46] Schlosser F, Maier BF, Jack O, Hinrichs D, Zachariae A, Brockmann D. COVID-19 lockdown induces disease-mitigating structural changes in mobility networks. Proceedings of the National Academy of Sciences. 2020.

[47] Environmental Protection Agency. What value of statistical life does EPA use? Washington, DC: Environmental Protection Agency, 2020. https://www.epa.gov/environmental-economics/mortality-risk-valuation#whatvalue (accessed October 27, 2020); 2020.

[48] Nurchis MC, Pascucci D, Sapienza M, Villani L, D'Ambrosio F, Castrini F, et al. Impact of the Burden of COVID-19 in Italy: Results of Disability-Adjusted Life Years (DALYs) and Productivity Loss. International Journal of Environmental Research and Public Health. 2020;17(12):4233.

[49] Ryan D, Toews C, Sanchirico JN, Armsworth PR. Implications of policy adjustment costs for fisheries management. Natural Resource Modeling. 2017;30(1):74–90.

[50] Edridge AW, Kaczorowska J, Hoste AC, Bakker M, Klein M, Loens K, et al. Seasonal coronavirus protective immunity is short-lasting. Nature medicine. 2020;26(11):1691–1693.

[51] Buckner JH, Chowell G, Springborn MR. Dynamic Prioritization of COVID-19 Vaccines When Social Distancing is Limited for Essential Workers. medRxiv. 2020.

[52] Ohmit SE, Thompson MG, Petrie JG, Thaker SN, Jackson ML, Belongia EA, et al. Influenza vaccine effectiveness in the 2011–2012 season: protection against each circulating virus and the effect of prior vaccination on estimates. Clinical infectious diseases. 2014;58(3):319–327.

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

[54] Rowthorn RE, Laxminarayan R, Gilligan CA. Optimal control of epidemics in metapopulations. Journal of the Royal Society Interface. 2009;6(41):1135–1144.

[55] Goldman SM, Lightwood J. Cost optimization in the SIS model of infectious disease with treatment. The BE Journal of Economic Analysis & Policy. 2002;2(1).

[56] Barrett S, Hoel M. Optimal disease eradication. Environment and Development Economics. 2007:627–652.

[57] Castonguay FM, Sokolow SH, De Leo GA, Sanchirico JN. Cost-effectiveness of combining drug and environmental treatments for environmentally transmitted diseases. Proceedings of the Royal Society B. 2020;287(1933):20200966.

[58] Castonguay FM, Lasserre P. L'exploitation de ressources naturelles non renouvelables en asymétrie d'information. L'Actualité économique. 2019;95(2-3).

[59] Von Bertalanffy L. A quantitative theory of organic growth (inquiries on growth laws. II). Human biology. 1938;10(2):181–213.

[60] New MB. Farming freshwater prawns: a manual for the culture of the giant river prawn (Macrobrachium rosenbergii). 428. Food & Agriculture Org.; 2002.

[61] Ranjeet K, Kurup B. Heterogeneous individual growth of Macrobrachium rosenbergii male morphotypes. Naga, The ICLARM Quarterly. 2002;25(2):13–18.

[62] Sampaio CM, Valenti WC. Growth curves for Macrobrachium rosenbergii in semi-intensive culture in Brazil. Journal of the World Aquaculture Society. 1996;27(3):353–358.

[63] Lalrinsanga P, Pillai BR, Patra G, Mohanty S, Naik NK, Sahu S. Length weight relationship and condition factor of giant freshwater prawn Macrobrachium rosenbergii (De Man, 1879) based on developmental stages, culture stages and sex. Turkish Journal of Fisheries and Aquatic Sciences. 2012;12(4):917–924.

[64] Nwosu F, Wolfi M. Population dynamics of the giant African river prawn Macrobrachium vollenhovenii Herklots 1857 (Crustacea, Palaemonidae) in the cross River estuary, Nigeria. West African Journal of Applied Ecology. 2006;9(1).

[65] Lorenzen K. The relationship between body weight and natural mortality in juvenile and adult fish: a comparison of natural ecosystems and aquaculture. Journal of fish biology. 1996;49(4):627–642.

[66] Sokolow SH, Lafferty KD, Kuris AM. Regulation of laboratory populations of snails (Biomphalaria and Bulinus spp.) by river prawns, Macrobrachium spp.(Decapoda, Palaemonidae): implications for control of schistosomiasis. Acta tropica. 2014;132:64–74.

[67] Halstead NT, Hoover CM, Arakala A, Civitello DJ, De Leo GA, Gambhir M, et al. Agrochemicals increase risk of human schistosomiasis by supporting higher densities of intermediate hosts. Nature communications. 2018;9(1):837.

[68] Sokolow SH, Huttinger E, Jouanard N, Hsieh MH, Lafferty KD, Kuris AM, et al. Reduced transmission of human schistosomiasis after restoration of a native river prawn that preys on the snail intermediate host. Proceedings of the National Academy of Sciences. 2015;112(31):9650–9655.

[69] Dasgupta S, Tidwell JH. A breakeven price analysis of four hypothetical freshwater prawn, Macrobrachium rosenbergii, farms using data from Kentucky. Journal of Applied Aquaculture. 2003;14(1-2):1–22.

[70] Miguel E, Kremer M. Worms: identifying impacts on education and health in the presence of treatment externalities. Econometrica. 2004;72(1):159–217.

[71] Ndamba J, Makaza N, Munjoma M, Gomo E, Kaondera K. The physical fitness and work performance of agricultural workers infected with Schistosoma mansoni in Zimbabwe. Annals of Tropical Medicine & Parasitology. 1993;87(6):553–561.

[72] Fenwick A, Figenschou B. The effect of Schistosoma mansoni infection on the productivity of cane cutters on a sugar estate in Tanzania. Bulletin of the World Health Organization. 1972;47(5):567.

[73] Kamel M, Moustafa Y, Foda N, Khashab S, Moemen M, Abo El Naga R. Impact of schistosomiasis on quality of life and productivity of workers. Eastern Mediterranean Health Journal. 2002;8(2-3):354–362.

[74] Weisbrod BA, Helminiak TW. Parasitic diseases and agricultural labor productivity. Economic Development and Cultural Change. 1977;25(3):505–522.

[75] Castonguay FM, Blackwood JC, Howerton E, Shea K, Sims C, Sanchirico JN. Spatial Allocation of Scarce COVID-19 Vaccines. medRxiv. 2021. Available from: `https://www.medrxiv.org/content/early/2021/03/14/2020.12.18.20248439`.