DISCUSSION PAPER SERIES

18271-1657977888

Epidemics in Space: Control, Targeting and Delegation

Darren Hoover and Flavio Toxvaerd

PUBLIC ECONOMICS

CEPR

ISSN 0265-8003

Epidemics in Space: Control, Targeting and Delegation

Darren Hoover and Flavio Toxvaerd

Discussion Paper 18271-1657977888
Published N/A
Submitted 16 July 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Public Economics

Epidemics in Space: Control, Targeting and Delegation

Abstract

We analyse optimal disease mitigation in a spatial model where disease spreads within and between interconnected regions. We characterise optimal strategies and emphasise the role of inter-regional coordination and policy targeting. Delegation of policy to regional planners achieves targeting without coordination, while a centrally determined uniform policy achieves coordination without targeting; both induce inefficiencies. For strongly connected regions, policy coordination is paramount, while for weakly connected regions, targeting becomes more important. Last, we analyse the value of reductions in integration, such as travel restrictions. We show that these may be non-monotone and sensitive to the underlying mitigation policy in place.

JEL Classification: C73, E61, I18, H75, R13

Keywords: Economic epidemiology, mitigation, spatial epidemiology, regional social planners, travel restrictions

Darren Hoover - dh559@cam.ac.uk
Faculty of Economics, University of Cambridge

Flavio Toxvaerd - fmot2@cam.ac.uk
Faculty of Economics, University of Cambridge and CEPR

Acknowledgements
We gratefully acknowledge constructive conversations on spatial economic models with Chryssi Giannitsarou and Tony Yates.

Epidemics in Space: Control, Targeting and Delegation*

Darren Hoover[1] and Flavio Toxvaerd†[2]

[1]Faculty of Economics, University of Cambridge
[2]Faculty of Economics, University of Cambridge and CEPR

Abstract

We analyse optimal disease mitigation in a spatial model where disease spreads within and between interconnected regions. We characterise optimal strategies and emphasise the role of inter-regional coordination and policy targeting. Delegation of policy to regional planners achieves targeting without coordination, while a centrally determined uniform policy achieves coordination without targeting; both induce inefficiencies. For strongly connected regions, policy coordination is paramount, while for weakly connected regions, targeting becomes more important. Last, we analyse the value of reductions in integration, such as travel restrictions. We show that these may be non-monotone and sensitive to the underlying mitigation policy in place.

Keywords: Economic epidemiology, mitigation, spatial epidemiology, regional social planners, travel restrictions.
JEL classification: C73, E61, I18, H75, R13.

*We gratefully acknowledge constructive conversations on spatial economic models with Chryssi Giannitsarou and Tony Yates.
†Corresponding author: fmot2@cam.ac.uk.

1 Introduction

As the COVID-19 pandemic has spread across the world, it has had an undeniable detrimental impact on health, prosperity and social well-being. One of the main methods used to control the spread of the disease has been the deployment of non-pharmaceutical interventions (NPIs). These measures involve altering human behaviour to reduce physical proximity and disease spread. While classical epidemiological analyses generally take human contacts and behaviour as exogenously given and constant over time, economic epidemiology instead models behaviour and disease dynamics in an integrated manner.[1] These studies typically consider a single isolated population and analyse the impact of mitigation and control measures on the spread of a disease over time. Diseases, however, spread not only over time, but also across space. This raises a number of important conceptual and practical questions about how to formulate and implement optimal public health measures. First, when a region's disease dynamics are intertwined with those of connected regions through travel or physical proximity, how does that influence the optimal mitigation policy? Second, what are the costs and benefits of coordinating such policies across regions? Third, how important is it to target mitigation policies to local conditions and what are the costs and benefits of non-targeted policies implemented across interdependent regions? Last, what is the value of reducing such interdependence, e.g. by the restriction of travel between regions or countries? We show that the answers to these questions depend on the degree of interconnectedness between jurisdictions in a straightforward way.

The role of space plays an intuitive role in the mechanics of disease spread. Proximally located individuals are more likely to physically interact and, thus, transmit an infection. On the other hand, the migration or commuting of infected individuals can allow for a disease to spread over large distances and be introduced to new regions. Yet, spatio-temporal dynamics has attracted surprisingly little attention in the fast-growing literature on economic epidemiology, especially that which seeks to provide practical policy advice. The relevance of the spatial aspect of disease spread, and how it conditions mitigation strategies across jurisdictions, was evident during the COVID-19 pandemic. Many countries allowed some inward and outward travel activity, while within-country travel restrictions were often imperfectly enforced, if they existed at all. Many interlinked countries have found themselves in different epidemiological conditions and have consequently enacted different mitigation policies (Yarmol-Matusiak et al., 2021). On a more local scale, some countries enacted policies at the state level. In the U.S., each state imposed restrictions of varying severity, while some states experienced higher prevalence

[1] E.g., Acemoglu et al. (2021); Rowthorn and Toxvaerd (2020); Makris (2021); Toxvaerd (2019, 2020).

of COVID-19 infections than others.[2,3] Even further, mitigation policies may instead be implemented at the county or city level. In England, following a national lockdown in November 2020, a region-specific policy scheme was implemented, in which different parts of the country were subject to different 'tiers' of COVID-19 restrictions.[4] As the scale becomes more local, individuals in neighbouring regions are more proximally located and have increased interaction. This entails a higher level of regional connection, through which a disease can more easily spread between regions.

In simple one-region homogeneous agent epidemic models, the issue of formulating optimal disease mitigation and control policies is now better understood, with a number of studies characterizing the use of lockdowns, treatment, vaccines and other measures in different disease contexts such as SI, SIS, SIRS and SIR specifications. One of the most robust results from this literature is that in equilibrium under decentralised decision making, i.e., under laissez-faire, there is typically suboptimal mitigation from a social perspective. As shown by Rowthorn and Toxvaerd (2020), the inefficiency of equilibrium stems from a combination of uninternalised externalities (the externality effect) and the fact that the decisions of non-atomistic individuals are conditioned by the fact that they cannot individually influence aggregate disease dynamics (the smallness effect).

Moving to a metapopulation setting, a number of new issues arise. For simplicity, consider a symmetric two-region SIR type model in which each region is connected to the other region by mutual travel flows. In such a model, we make two central observations that will help clarify our analysis and findings. Consider a situation in which one of the regions is seeded with infections from an external source, while the other region is initially fully susceptible. Our first observation is that because of the asymmetric seeding, the subsequent disease dynamics will typically be imperfectly synchronised. The extent to which the disease paths move in parallel, all else equal, depends on how strongly interconnected the regions are. The stronger the coupling is, the more similar the disease dynamics (Keeling and Rohani, 2008). The second observation is that when there are multiple interconnected regions, there are necessarily policy spillovers in the sense that any mitigation implemented in one region will impact disease dynamics, welfare and mitigation incentives in the other region. But note also that the importance of these two findings, namely (i) synchrony and (ii) spillovers depends on the strength of interdependence between the regions. When regions are only weakly coupled, disease dynamics become more asymmetric but policy spillovers become weaker. Similarly, as the coupling of regions becomes stronger, disease dynamics become more symmetric but policy spillovers become stronger. This will turn out to condition policy in important ways.

These two observations are reflected in socially optimal mitigation policies. Consider

[2]https://www.cleveland.com/metro/2020/03/50-states-of-coronavirus-how-every-state-in-the-us-has-responded-to-the-pandemic.html

[3]https://www.nytimes.com/interactive/2021/us/covid-cases.html

[4]https://www.bbc.co.uk/news/uk-55078888

a global social planner tasked with choosing economically costly mitigation policies in the two regions with a view to maximise overall (global) social well-being. The optimal policy chosen by such a social planner will have two central properties, namely that it is (i) coordinated and (ii) targeted. Coordination simply means that the mitigation policy implemented in either region is sensitive to the impact that it has on the disease dynamics and welfare in the other region and vice versa. In other words, the optimal policy internalises the inter-regional externalities of mitigation. Targeting simply means the policy that is implemented is sensitive to regional differences in disease dynamics. As the epidemic progresses differently in the two (interconnected) regions, optimal policy will reflect those differences.

To tease out how those considerations depend on the interconnectedness of the regions, we consider two departures from the globally socially optimal policy. First, we consider the optimal policy chosen by a global social planner who is restricted to a uniform policy, i.e., a mitigation path that is not targeted to the different regions. Note that this policy is coordinated in the sense that the planner's objective is still to maximise overall social welfare and so is considering the effects that the policy in one region has on the welfare of the other region. Second, we consider the policy that emerges when mitigation policies are delegated to regional social planners who seek to maximise social welfare in each their own region. In this case, the resulting policies are no longer coordinated (in fact, the two regional planners are now playing a non-cooperative game and the resulting policies are the equilibrium of this game), but the policies are still targeted, because they reflect the different disease dynamics prevailing in the two regions.

The two alternative policy approaches, the uniform policy (non-targeted but coordinated) and the delegated policy (targeted but uncoordinated), are both deviations from the global social optimum and reflect different inefficiencies caused by ignoring the two properties of the first-best policy. But these inefficiencies stem from the synchrony and spillover effects discussed earlier and, as such, vary in strength with the strength of inter-regional coupling. For weakly coupled regions, disease dynamics are relatively asymmetric and thus an optimal targeted policy will induce different mitigation paths in different regions. In addition, in this setting, policy spillovers will be weak and there is less need for coordination. The upshot of this is that when regions are weakly connected, delegation of policy to regional social planners performs relatively well, whereas a global social planner restricted to uniform mitigation policies would perform poorly. Conversely, with strongly coupled regions, disease dynamics are very similar in the two regions, while policy spillovers are strong. Optimal policy will in this case induce similar mitigation paths in both regions and the main concern becomes one of internalising the inter-regional effects of regional mitigation policies. In this scenario, a uniform mitigation policy will perform relatively well, whereas the outcome of delegating policy to regional policy makers will perform relatively poorly, because they fail to account for inter-regional externalities.

In terms of public health policy, our findings suggest that, in general, non-targeted policies applied to similar and well-integrated neighbouring regions can go far in achieving desirable outcomes. In contrast, far-flung non-contiguous regions may benefit more from delegated and therefore targeted policies, even if these are not coordinated with the policies implemented in other regions.

Finally, we analyse the benefits of travel restrictions, interpreted as a reduction in cross-region interdependence. We do so by computing the regional and total willingness to pay for a reduction in the cross-region interaction rate. Including endogenous mitigation in the model results in a non-monotonicity of the willingness to pay, leading to notable differences in the benefits of travel restrictions. In some cases, these differences are substantial enough to result in net impacts of opposite signs. Furthermore, travel restrictions are more likely to be beneficial when mitigation policies are set by uncoordinated regional social planners.

Related literature

Due to the COVID-19 pandemic, work involving the economic analysis of mitigation in singular populations has been expanding rapidly. Relevant early work includes that by Sethi (1978), Reluga (2010), Chen et al. (2011), Chen (2012) and Fenichel (2013), which assesses socially optimal and equilibrium mitigation levels in SIS and SIR models. More recently, Toxvaerd (2019, 2020) analyses equilibrium and socially optimal mitigation levels when individuals are forward-looking with perfect foresight. Similarly, Makris and Toxvaerd (2020) explore equilibrium and optimal mitigation levels when a pharmaceutical intervention is expected at a certain date. Makris (2021) and Acemoglu et al. (2021) incorporate population heterogeneity into the SIR model, which they use to discuss equilibrium and socially optimal mitigation levels in relation to the COVID-19 pandemic.

For work in economic epidemiology with spatial models, Rowthorn et al. (2009) and Ndeffo Mbah and Gilligan (2011) study optimal pharmaceutical treatment allocation strategies for SIS and SIRS metapopulation models. Birge et al. (2022) and Fajgelbaum et al. (2021) use commuter metapopulation models of an SEIR disease with asymptomatic infection to assess socially optimal targeted lockdown policies against COVID-19. Acemoglu et al. (2020) consider a network model with an SIR type disease framework and characterise equilibrium mitigation levels and corresponding socially optimal testing policies. Bognanni et al. (2020) propose an SIR model in which myopic individuals choose to mitigate. After introducing a metapopulation framework and calibrating the parameters to the COVID-19 pandemic, they analyse the equilibrium and socially optimal levels of mitigation.

In contrast to the existing literature, we propose a spatial epidemiological model in

[4]https://www.nytimes.com/2020/10/19/world/europe/coronavirus-manchester-britain.html

which mitigation is determined both endogenously and dynamically. This framework involves individuals who are not myopic, but rather forward-looking when making their mitigation decision. Furthermore, to the best of our knowledge, this paper is the first to analyse the costs of uncoordinated regional social planners and the benefits of travel restrictions within an epidemiological framework of endogenous mitigation.

This paper is structured as follows. In Section 2, we outline the main components of a metapopulation model with mitigation. We also provide a discussion of the disease dynamics in the epidemiological benchmark where no mitigation is present. In Section 3, we discuss the determination of equilibrium mitigation under decentralised decision making, along with the impact that one region's mitigation levels have on a neighbouring region's best response. Section 4 first considers the socially optimal mitigation levels of a global social planner. Then, we consider the case of mitigation under uncoordinated regional social planners and the (constrained optimal) uniform policy. Section 5 contains an extension in which we analyse the benefits of travel restrictions under the different decision-making settings. Lastly, we present our conclusions in Section 6.

2 Model

We first set out the model of uncontrolled epidemic dynamics and subsequently introduce the economic model of mitigation. Last, we briefly explain the dynamics of the disease without mitigation.

2.1 Spatial SIR model

Consider a metapopulation SIR model of two coupled regions $j = 1, 2$ with homogeneous continuum populations, each normalised to size $N = 1$. Time is continuous and at each moment $t \geq 0$, individuals in each region j can belong to either the set of susceptible, infected, or recovered persons, represented by $\mathcal{S}_j(t)$, $\mathcal{I}_j(t)$, and $\mathcal{R}_j(t)$ respectively. The proportion of individuals in region j at time t in each disease compartment is then represented by $S_j(t)$, $I_j(t)$, and $R_j(t)$, where $S_j(t), I_j(t), R_j(t) \in [0, 1]$ and $S_j(t) + I_j(t) + R_j(t) = 1$. Thus, temporarily ignoring the prospect of mitigation, the two-region SIR disease dynamics are given by

$$\dot{S}_j(t) = -\beta S_j(t) I_j(t) - \varepsilon\beta S_j(t) I_{-j}(t) \tag{2.1a}$$

$$\dot{I}_j(t) = \beta S_j(t) I_j(t) + \varepsilon\beta S_j(t) I_{-j}(t) - \gamma I_j(t) \tag{2.1b}$$

$$\dot{R}_j(t) = \gamma I_j(t) \tag{2.1c}$$

where $j \in \{1, 2\}$, $\beta > 0$ is the transmission rate, $\gamma > 0$ is the recovery rate, and $\varepsilon \in [0, 1]$ captures the reduced force of transmission from infected individuals of the neighbouring

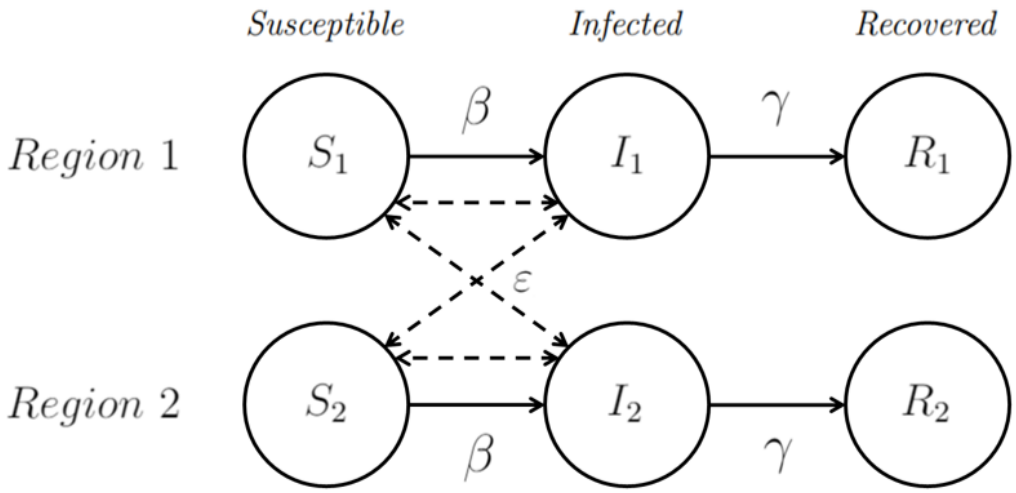

Figure 1: Flow diagram of disease compartment transitions in the metapopulation SIR model. Solid lines denote transition flows between disease compartments. Dashed lines denote disease-transmitting contact structure.

region. Figure 1 represents the mechanics of disease state transitions in this metapopulation SIR model (2.1).[5]

Formulation (2.1) assumes that the cross-region transmission rate is symmetrical for both regions. Alternatively, one could consider this cross-region interaction to be asymmetric, such that region 1 has a stronger impact on region 2 than 2 has on 1, or vice versa. However, such a specification would be functionally equivalent to the case where one region has more infected individuals than the other.

2.2 Mitigation

Each individual receives at every moment t $ge0$ a health-related utility flow π_τ corresponding to their infection state τ at the time. We assume that being susceptible or recovered is preferred to being infected. Furthermore, we allow for the possibility of long-term adverse effects from undergoing infection, that extend beyond the moment of recovery. Thus, $\pi_{\mathcal{S}} \geq \pi_{\mathcal{R}} \geq \pi_{\mathcal{I}}$. Additionally, each individual i in region j is able to choose their mitigation level $m_j^i(t) \in [0, 1]$ at each moment in time. This mitigation level can be interpreted, for example, as the proportion of an individual's interactions that they choose to forego. Choosing to mitigate comes at a personal cost $\frac{c}{2} m_j^i(t)^2$. However, doing so reduces the number of contacts the individual makes with the rest of the population, thus reducing their chance of engaging in a disease-transmitting interaction. Altogether, the

[5]Note that the model can be extended to explicitly include an additional compartment for disease-induced death. We forego this extension primarily for the sake of expositional simplicity. Furthermore, particularly for pseudo mass-action transmission mechanics or low mortality rates, incorporating disease-induced deaths into the model is effectively equivalent to the deceased individuals instead moving to the recovered compartment (Keeling and Rohani, 2008).

instantaneous net utility earned by individual i in infection compartment τ_j at time $t \geq 0$ is written as

$$u_i(t) = \pi_\tau - \frac{c}{2} m_j^i(t)^2 \tag{2.2}$$

Note that, due to the quadratic form of the mitigation costs, the instantaneous utility function is concave in the individual's choice variable $m_j^i(t) \in [0, 1]$. Throughout, the future will be discounted at a rate $\rho > 0$.

We assume individuals to be self-serving and only interested in maximising their own payoff. As a result, infected and recovered individuals will never mitigate. This is because any measure of mitigation $m_j^i(t) > 0$ is costly and only serves to reduce the risk of getting infected, which these individuals do not benefit from. Thus, the decision problem of choosing to mitigate is solely faced by susceptible individuals.

The aggregate mitigation efforts undertaken by the susceptible populations impact the disease dynamics in each region by reducing the number of disease-transmitting interactions. Let $\bar{m}_j(t)$ denote the average mitigation level of susceptible individuals in region j, calculated as

$$\bar{m}_j(t) \equiv \frac{\sum_{i \in \mathcal{S}_j(t)} m_j^i(t)}{S_j(t)} \tag{2.3}$$

Using $\bar{m}_j(t)$, the metapopulation SIR model with mitigation is described by the equations

$$\dot{S}_j(t) = -\left(1 - \bar{m}_j(t)\right) \beta S_j(t) \left(I_j(t) + \varepsilon I_{-j}(t)\right) \tag{2.4a}$$

$$\dot{I}_j(t) = \left(1 - \bar{m}_j(t)\right) \beta S_j(t) \left(I_j(t) + \varepsilon I_{-j}(t)\right) - \gamma I_j(t) \tag{2.4b}$$

$$\dot{R}_j(t) = \gamma I_j(t) \tag{2.4c}$$

for $j = 1, 2$.

2.3 Uncontrolled disease dynamics

To set the stage for subsequent analysis, we first discuss the epidemiological benchmark in which no individuals mitigate. This will act as a frame of reference when discussing outcomes under endogenous mitigation.

As the SIR model is not analytically tractable, numerical methods and simulations are the main tools for gaining insight into the model's disease dynamics.[6] However, there are some basic epidemiological principles which can be derived from such models. One of these principles, the 'threshold phenomenon', states that a disease can only invade and spread in a population if its *basic reproductive ratio* (R_0) is greater than 1 (Keeling and Rohani, 2008). This basic reproductive ratio, which is one of the most important measures used in epidemiological models, represents the average number of secondary

[6]Further details on the numerical solution methodology are provided in Appendix C.

infections caused by an initial infection in an entirely susceptible population. In the case of the symmetric metapopulation SIR model (2.1), we have

$$R_0 \equiv \frac{\beta\,(1+\varepsilon)}{\gamma} \tag{2.5}$$

If $R_0 \leq 1$, then any initial infection in a population will die out and fail to take hold. In other words, an epidemic will not take place. To ensure initial infections are 'successful' (at least in the absence of mitigation) and to allow for more interesting analysis, we impose the following assumption:

Assumption 1. $\beta\,(1+\varepsilon) > \gamma$.

Throughout, we consider two specifications of the model: one where ε is low and one where ε is high. This is to see how the results depend on the level to which regions are inter-connected. In order to keep the two model specifications comparable in terms of disease severity, the transmission rate parameter β is adjusted to maintain the same basic reproductive ratio R_0. The two model specifications are otherwise identical in all parameter values.[7] Furthermore, in each simulation, a small number of infections are initially introduced into only region $j = 1$.

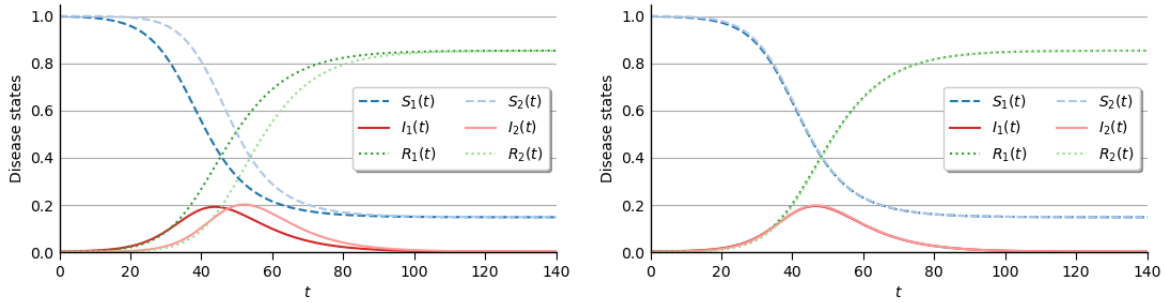

(a) Low ε. (b) High ε.

Figure 2: Disease dynamics without mitigation.

As illustrated in Figures 2 and 3, both the weakly and strongly coupled model specifications demonstrate characteristics that are common in SIR-type models. Infection prevalence within each region is hump-shaped over time. During the initial growth phase of the epidemic, infection prevalence grows more and more rapidly. During the peak phase, infection prevalence stabilises and then starts to diminish. Infection levels continue to dissipate as the epidemic ends. In the long-run, infections disappear ($I_j(t) \to 0$) and the disease states are stable ($\dot{S}_j(t) \approx \dot{I}_j(t) \approx \dot{R}_j(t) \approx 0$). The hump-shape pattern is

[7]For the disease dynamics parameters, we use $R_0 = 2.25$ as this is similar to estimates made for the COVID-19 pandemic (Kucharski et al., 2020; Wu et al., 2020). Additionally, we use a recovery rate of $\gamma = 1/7.5$, which implies an average infectious period of 7.5 days. This is also close to some estimates made for the COVID-19 pandemic (Byrne et al., 2020). The other parameter values used in these simulations are: $I_1(0) = 0.001,\ I_2(0) = 0,\ R_1(0) = R_2(0) = 0,\ (a)$ low $\varepsilon = 0.025$ (b) high $\varepsilon = 0.25$.

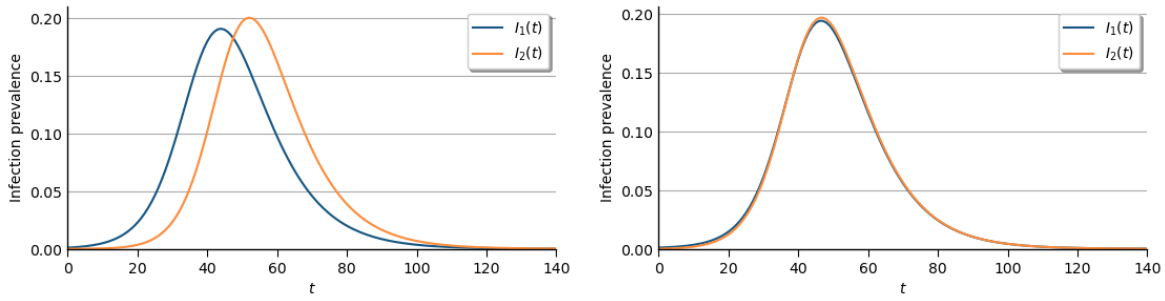

(a) Low ε.

(b) High ε.

Figure 3: Infection prevalence without mitigation.

also present in a region's *infection force*, which is the rate at which a susceptible individual incurs an infection. In the symmetric metapopulation model, the infection force in region j is $\beta\left(I_j(t) + \varepsilon I_{-j}(t)\right)$.

Some qualitative differences can be noted between the two model specifications. In the case of strongly coupled regions, the initial injection of infections in region 1 immediately spreads to region 2. Throughout the epidemic, the two regions are essentially symmetric, as their disease states hold almost identical values. This is not so in the case of weakly coupled regions. In this specification, it takes time for the infection to spread to the second region. This asymmetry is maintained throughout the epidemic, as there is a delay between region 1 and region 2's infection peak. As can be seen in Figure 3, the peak infection prevalence is slightly higher in region 2. This is due to the fact that as the infection is initially introduced in region 1, the initial level of susceptibles is higher in region 2 (i.e., $S_2(0) > S_1(0)$), allowing for a larger epidemic in the region (Keeling and Rohani, 2008). These asymmetries will turn out to have interesting effects once disease mitigation is introduced in the different regions.

3 Laissez-faire

In this section, we consider the case of mitigation under decentralised decision making, i.e., laissez-fare, in which every individual independently and non-cooperatively chooses how much to mitigate. In this setting, each individual i in region j is forward-looking with perfect foresight, and has the goal of choosing their mitigation level $m_j^i(t)$ at each moment in order to maximise their expected discounted utility. Infected and recovered individuals will always choose not to mitigate. Let $p^i_{\tau_j}(t)$ denote the probability of an initially susceptible individual i being in infection state τ_j at time $t \geq 0$. The susceptible

individual's decision problem can then be written as

$$\begin{aligned}
\max_{m_j^i(t)\in[0,1]} \int_0^\infty e^{-\rho t} \left[p_{S_j}^i(t) \left(\pi_{\mathcal{S}} - \frac{c}{2} m_j^i(t)^2 \right) + p_{I_j}^i(t)\pi_{\mathcal{I}} + p_{R_j}^i(t)\pi_{\mathcal{R}} \right] dt \\
s.t. \quad \dot{p}_{S_j}^i(t) &= -\left(1 - m_j^i(t)\right) \beta \left(I_j(t) + \varepsilon I_{-j}(t)\right) p_{S_j}^i(t) \\
\dot{p}_{I_j}^i(t) &= \left(1 - m_j^i(t)\right) \beta \left(I_j(t) + \varepsilon I_{-j}(t)\right) p_{S_j}^i(t) - \gamma p_{I_j}^i(t) \\
\dot{p}_{R_j}^i(t) &= \gamma p_{I_j}^i(t) \\
p_{S_j}^i(0) &= 1\,, \quad p_{I_j}^i(0) = p_{R_j}^i(0) = 0
\end{aligned} \tag{3.1}$$

Since each individual is non-atomistic, their choice of mitigation level $m_j^i(t)$ has a negligible impact on the population's average mitigation level $\bar{m}_j(t)$ and thus aggregate dynamics. Thus, the disease dynamics and infection prevalences are taken as exogenously given in the individual's optimisation problem.

The decision problem (3.1) of an individual i in region j is equivalent to maximising the current-value Hamiltonian:[8]

$$\begin{aligned}
H_j^i = {} & p_{S_j}^i(t) \left(\pi_{\mathcal{S}} - \frac{c}{2} m_j^i(t)^2 \right) + p_{I_j}^i(t)\pi_{\mathcal{I}} + p_{R_j}^i(t)\pi_{\mathcal{R}} \\
& - \lambda_{S_j}(t) \left(1 - m_j^i(t)\right) \beta \left(I_j(t) + \varepsilon I_{-j}(t)\right) p_{S_j}^i(t) \\
& + \lambda_{I_j}(t) \left[\left(1 - m_j^i(t)\right) \beta \left(I_j(t) + \varepsilon I_{-j}(t)\right) p_{S_j}^i(t) - \gamma p_{I_j}^i(t) \right] \\
& + \lambda_{R_j}(t) \gamma p_{I_j}^i(t)
\end{aligned} \tag{3.2}$$

in which $\lambda_{\tau_j}(t)$ is the current-value costate variable of disease compartment τ_j.[9] As usual, the current-value costate variable's laws of motion is given by (Caputo, 2005)

$$\dot{\lambda}_{\tau_j}(t) = \rho \lambda_{\tau_j}(t) - \frac{\partial H_j^i}{\partial p_{\tau_j}^i(t)} \tag{3.3}$$

For the costate variables of the laissez-faire decision problem, this yields

$$\begin{aligned}
\dot{\lambda}_{S_j}(t) = {} & \lambda_{S_j}(t)\rho + \left(1 - m_j^i(t)\right) \beta \left(I_j(t) + \varepsilon I_{-j}(t)\right) \left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right) \\
& - \pi_{\mathcal{S}} + \frac{c}{2} m_j^i(t)^2
\end{aligned} \tag{3.4a}$$

$$\dot{\lambda}_{I_j}(t) = \lambda_{I_j}(t) \left(\rho + \gamma\right) - \lambda_{R_j}\gamma - \pi_{\mathcal{I}} \tag{3.4b}$$

$$\dot{\lambda}_{R_j}(t) = \lambda_{R_j}(t)\rho - \pi_{\mathcal{R}} \tag{3.4c}$$

Differentiating (3.2) with respect to $m_j^i(t)$ yields the necessary condition for maximi-

[8] This statement, along with the rest of the analysis in this paper, assumes an interior solution. If the mitigation level that maximises the current-value Hamiltonian does not lie between 0 and 1 at any time $t \in [0, \infty)$, the current-value Hamiltonian would need to be altered to include the relevant boundary constraints (Seierstad and Sydsaeter, 1987).

[9] See Appendix A for discussion on the economic interpretation and transversality conditions of the current-value costate variables.

sation

$$\frac{\partial H_j^i}{\partial m_j^i(t)} = -p_{S_j}^i(t) c m_j^i(t) + \left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right) \beta \left(I_j(t) + \varepsilon I_{-j}(t)\right) p_{S_j}^i(t) = 0 \tag{3.5}$$

which, assuming $p_{S_j}^i(t) > 0$, gives the condition[10]

$$m_j^i(t) = \frac{\beta \left(I_j(t) + \varepsilon I_{-j}(t)\right) \left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right)}{c} \tag{3.6}$$

Let $m_j^{i*}(t)$ denote the equilibrium level of mitigation taken by a susceptible individual i in region j, being the solution to the necessary maximisation condition (3.6) subject to $m_j^i(t) \in [0, 1]$ and the costate variables' laws of motion (3.4). Since all susceptible individuals in a given region are assumed to be identical, they all face the same decision problem and therefore have the same solution for $m_j^{i*}(t)$. Thus, we define $m_j^*(t) \equiv m_j^{i*}(t), \ \forall i \in \mathcal{S}_j(t)$ as the mitigation level undertaken by all susceptible individuals in region j. The disease dynamics in the laissez-faire setting are hence described by the metapopulation SIR model with mitigation (2.4) in which $\bar{m}_j(t) = m_j^*(t)$, for $j = 1, 2$. These disease dynamics are illustrated in Figure 4.

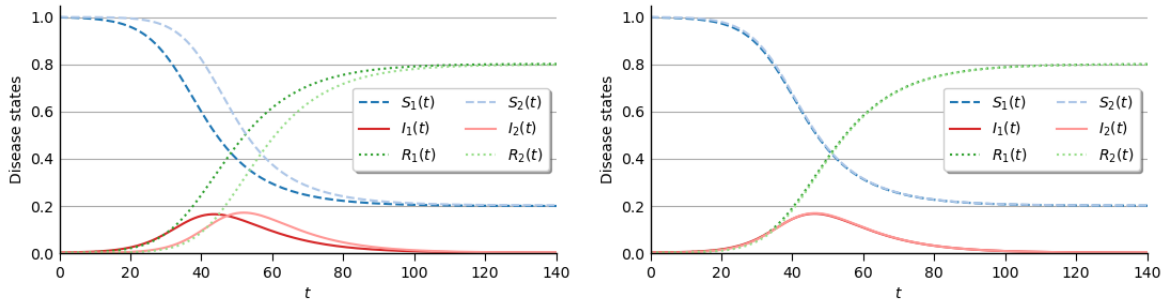

(a) Low ε. (b) High ε.

Figure 4: Disease dynamics under laissez-faire mitigation.

As made clear in Figure 5, including endogenous mitigation under laissez-faire results in a flattened infection prevalence. Infection levels during the peak of the epidemic are lowered. However, the duration of the epidemic is extended, as infection prevalence is higher towards the end of the epidemic. Overall, cumulative infections in all regions are lower under laissez-faire mitigation than in the epidemiological benchmark (as $R_j(\infty)$ is lower). Furthermore, the total welfare within each region is higher in the laissez-faire setting.

Although the duration of the epidemic is extended when including mitigation, there is no noticeable impact in these simulations on the timing of the epidemic peak, nor the speed at which infections spread to the second region. This is due to the fact that

[10] As the decision problem is faced solely by a susceptible individual, it is only relevant when $p_{S_j}^i(t) > 0$.

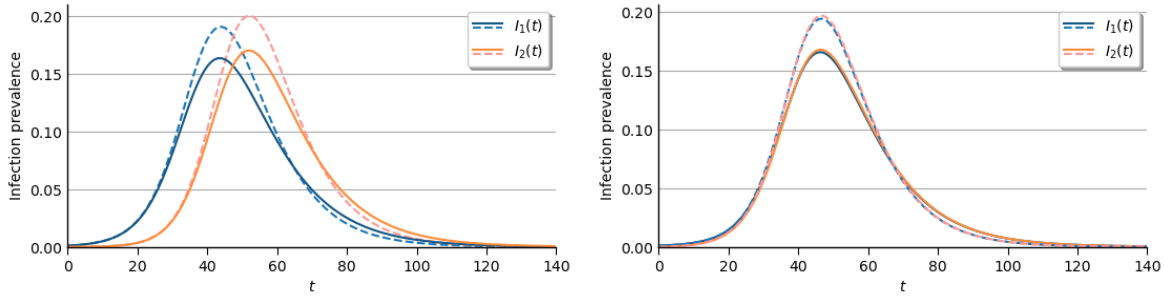

(a) Low ε. (b) High ε.

Figure 5: Infection prevalence under laissez-faire mitigation (solid lines) and in the benchmark without mitigation (dashed lines).

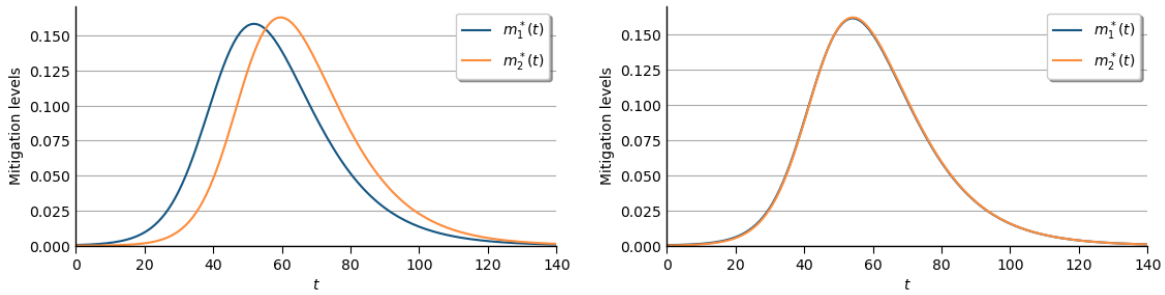

(a) Low ε. (b) High ε.

Figure 6: Mitigation levels under laissez-faire mitigation.

infection prevalence (and, thus, risk of disease transmission) is low in the early stages of the epidemic, resulting in correspondingly low mitigation levels. Once infection levels increase and the disease starts to spread more rapidly, mitigation levels rise as well. It is only then that mitigation's impact on disease dynamics becomes noticeable.

As can be seen in Figure 6, in both the case of weakly and strongly connected regions, the equilibrium mitigation levels in a given region exhibits a hump-shaped pattern, with a peak occurring after the region's peak infection prevalence (and, correspondingly, their peak infection force). This is because, as seen in (3.6), the mitigation level $m_j^*(t)$ at a given time t depends on the region's infection force $\beta\left(I_j(t) + \varepsilon I_{-j}(t)\right)$ and the value of avoiding infection $\left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right)$. This latter term is positive (as we assume infection is undesired) and increasing throughout the course of the epidemic. The reason for its increasing nature is that if one avoids infection at the beginning, they still face a large risk of getting infected later. However, if one manages to stay uninfected towards the end of the epidemic, then they are more likely to exit the epidemic without ever becoming infected. Remaining uninfected therefore becomes more valuable as the epidemic progresses and the chance of a delayed future infection decreases. Thus, the equilibrium mitigation levels mirror the region's infection force curve, but occurs at a later time. These qualitative patterns are similarly found in previous studies using single-population SIR models, such as Makris and Toxvaerd (2020), Toxvaerd (2020) and Farboodi et al. (2021).

3.1 Mitigation's inter-regional strategic nature

One of the immediate questions that arises when considering endogenous mitigation in a spatial model is what the inter-regional strategic effects are. In other words, how do mitigation levels in one region impact the decision to mitigate in another region. To assess this question, we consider a thought experiment in which individuals in region 1 are given the same option to mitigate as before, but individuals in region 2 are not given the opportunity to mitigate. As $m^i_{-j}(t)$ does not appear in the individual's decision problem (3.1) and the disease dynamics are taken as exogenously given, the optimality condition (3.6) and costate variables' laws of motion (3.4) still hold for susceptible individuals in region 1. Let $\tilde{m}^*_1(t)$ denote the equilibrium mitigation level of susceptible individuals in region 1. The disease dynamics under this alternative setting, which we will henceforth refer to as the 'quasi-decentralised' decision-making setting, are then described by the metapopulation SIR model with mitigation (2.4) where $\bar{m}_1(t) = \tilde{m}^*_1(t)$ and $\bar{m}_2(t) = 0$, and are illustrated in Figure 7.

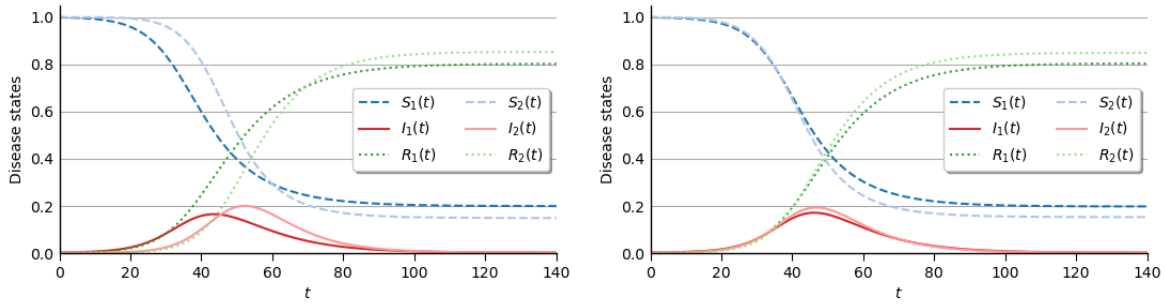

(a) Low ε. (b) High ε.

Figure 7: Disease dynamics under the quasi-decentralised setting.

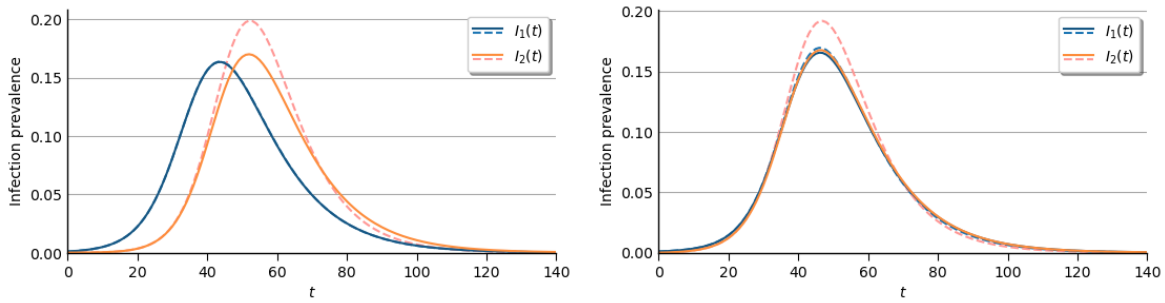

(a) Low ε. (b) High ε.

Figure 8: Infection prevalence in the laissez-faire setting (solid lines) and the quasi-decentralised setting (dashed lines).

Figure 8 shows that, relative to the quasi-decentralised setting, the laissez-faire setting (in which mitigation takes place in both regions) results in a flattening of region 2's infection prevalence curve, and a marginal flattening of region 1's infection prevalence

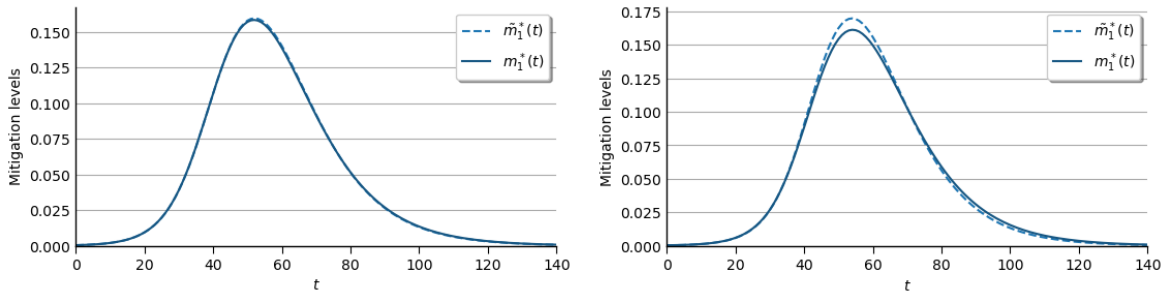

(a) Low ε. (b) High ε.

Figure 9: Mitigation levels in region 1 in the laissez-faire setting (solid lines) and the quasi-decentralised setting (dashed lines).

curve. This flattening of infection prevalences results in similar changes to the infection force experienced in region 1: lower values during the epidemic peak, but higher values towards the end of the epidemic. The magnitude of these changes are also greater with higher values of ε.

Region 2's mitigation levels, as they do not appear in the objective function of susceptible individuals in region 1, do not have a direct impact on mitigation levels in region 1. Rather, they indirectly impact 1's mitigation levels through altering the disease dynamics. As the decision to mitigate in region 2 results in a flattening of the infection force in region 1, there is a corresponding impact on the mitigation levels in region 1. This is demonstrated in Figure 9, in which mitigation in region 2 induces individuals in region 1 to do less mitigation during the epidemic peak, but more mitigation towards the end. The extent to which mitigation in region 1 changes is increasing with higher values of ε. Thus, the inter-regional strategic impact of mitigation is not constant over time. Higher mitigation levels in the neighbouring region induce lower mitigation levels at home (strategic substitutes) during the growth and peak phases of the epidemic, but induce higher mitigation levels at home (strategic complements) towards the end of the pandemic.[11]

Compared to the laissez-faire setting, the quasi-decentralised setting results in more cumulative infections and less social welfare in region 2. Additionally, there are marginally more cumulative infections and marginally less social welfare in region 1. The inter-regional impact on the cumulative infections and welfare in region 1, while still being small compared to the intra-regional impact, is higher in the case of strongly connected regions. This suggests that an individual's decision to mitigate has a positive inter-regional externality, which is larger in magnitude as ε increases.

[11]These results are robust to changing the region of initial infection. Furthermore, similar results are found when using different approaches to altering region 2's mitigation levels. For example, Appendix B considers the impact of increasing a constant mitigation level in region 2.

4 Centralised decision making

In this section, we consider mitigation under centralised decision making, in which a social planner aims to maximise overall social welfare. An individual's decision to mitigate exerts a positive externality on other individuals. By mitigating, one reduces their chance of becoming infected, which subsequently reduces their chance of infecting others in the future, both within their own region (*intra-regional* infection) and in the neighbouring region (*inter-regional* infection). Mitigation's externalities can correspondingly be broken into two components: the intra-regional externalities and the inter-regional externalities. If the susceptible individual's decision to mitigate were instead centralised, these intra-regional and inter-regional externalities could be accounted for. We therefore consider two different spatial aggregations of centralised decision making. Namely, where there is *global* decision making and *regional* decision making.

4.1 Global social planner

Under centralised decision making, we first consider the scenario where the decision to mitigate is fully coordinated between the two regions through a *global* social planner. In this setting, the global social planner is able to set the level of mitigation taken by susceptible individuals in each location with the goal of maximising aggregate social welfare across all regions. Unlike in the laissez-faire setting, the global social planner is able to set all susceptible individuals' mitigation efforts in both locations, and is thus able to directly control the disease dynamics. The global social planner takes this into account in their decision, and no longer takes the disease dynamics as given. Denoting the mitigation level chosen by the social planner for susceptible individuals in location j at time t as $m_j(t)$, the targeted global social planner's decision problem is written as

$$\begin{aligned}\max_{m_1(t),m_2(t)\in[0,1]} \int_0^\infty e^{-\rho t}\Big[&S_1(t)\left(\pi_{\mathcal{S}} - \frac{c}{2}m_1(t)^2\right) + I_1(t)\pi_{\mathcal{I}} + R_1(t)\pi_{\mathcal{R}} \\ &+ S_2(t)\left(\pi_{\mathcal{S}} - \frac{c}{2}m_2(t)^2\right) + I_2(t)\pi_{\mathcal{I}} + R_2(t)\pi_{\mathcal{R}}\Big]dt\end{aligned} \tag{4.1}$$

subject to, for $j = 1, 2$,

$$\dot{S}_j(t) = -\left(1 - m_j(t)\right)\beta S_j(t)\left(I_j(t) + \varepsilon I_{-j}(t)\right) \tag{4.2a}$$

$$\dot{I}_j(t) = \left(1 - m_j(t)\right)\beta S_j(t)\left(I_j(t) + \varepsilon I_{-j}(t)\right) - \gamma I_j(t) \tag{4.2b}$$

$$\dot{R}_j(t) = \gamma I_j(t) \tag{4.2c}$$

Denoting the costate variable of disease compartment τ_j as $\lambda_{\tau_j}(t)$, the global social

planner's current-value Hamiltonian is

$$\begin{aligned} H = {} & S_1(t)\left(\pi_{\mathcal{S}} - \frac{c}{2}m_1(t)^2\right) + I_1(t)\pi_{\mathcal{I}} + R_1(t)\pi_{\mathcal{R}} \\ & + S_2(t)\left(\pi_{\mathcal{S}} - \frac{c}{2}m_2(t)^2\right) + I_2(t)\pi_{\mathcal{I}} + R_2(t)\pi_{\mathcal{R}} \\ & - \lambda_{S_1}(t)\left(1 - m_1(t)\right)\beta S_1(t)\left(I_1(t) + \varepsilon I_2(t)\right) \\ & - \lambda_{S_2}(t)\left(1 - m_2(t)\right)\beta S_2(t)\left(I_2(t) + \varepsilon I_1(t)\right) \\ & + \lambda_{I_1}(t)\left[\left(1 - m_1(t)\right)\beta S_1(t)\left(I_1(t) + \varepsilon I_2(t)\right) - \gamma I_1(t)\right] \\ & + \lambda_{I_2}(t)\left[\left(1 - m_2(t)\right)\beta S_2(t)\left(I_2(t) + \varepsilon I_1(t)\right) - \gamma I_2(t)\right] \\ & + \lambda_{R_1}(t)\gamma I_1(t) \\ & + \lambda_{R_2}(t)\gamma I_2(t) \end{aligned} \tag{4.3}$$

As before, the laws of motion for the current-value costate variables are derived as

$$\dot{\lambda}_{\tau_j}(t) = \rho\lambda_{\tau_j}(t) - \frac{\partial H}{\partial \tau_j(t)} \tag{4.4}$$

which yields, for $j = 1, 2$:

$$\begin{aligned} \dot{\lambda}_{S_j}(t) = {} & \lambda_{S_j}(t)\rho + \left(1 - m_j(t)\right)\beta\left(I_j(t) + \varepsilon I_{-j}(t)\right)\left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right) \\ & - \pi_{\mathcal{S}} + \frac{c}{2}m_j(t)^2 \end{aligned} \tag{4.5a}$$

$$\begin{aligned} \dot{\lambda}_{I_j}(t) = {} & \lambda_{I_j}(t)\left(\rho + \gamma\right) + \left(1 - m_j(t)\right)\beta S_j(t)\left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right) \\ & + \left(1 - m_{-j}(t)\right)\varepsilon\beta S_{-j}(t)\left(\lambda_{S_{-j}}(t) - \lambda_{I_{-j}}(t)\right) - \lambda_{R_j}(t)\gamma - \pi_{\mathcal{I}} \end{aligned} \tag{4.5b}$$

$$\dot{\lambda}_{R_j}(t) = \lambda_{R_j}(t)\rho - \pi_{\mathcal{R}} \tag{4.5c}$$

Differentiating the current-value Hamiltonian (4.3) with respect to $m_j(t)$ and assuming $S_j(t) > 0$ leads to the necessary conditions for optimality

$$m_j(t) = \frac{\beta\left(I_j(t) + \varepsilon I_{-j}(t)\right)\left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right)}{c}, \quad j = 1, 2 \tag{4.6}$$

We denote the socially optimal mitigation levels in region j under the targeted global social planner setting as $m_j^G(t)$, being the solutions to the necessary optimality conditions (4.6), subject to $m_j(t) \in [0, 1]$ and the costate variables' laws of motion (4.5). The resulting SIR disease dynamics for each region are then given by the metapopulation SIR model with mitigation (2.4) where $\bar{m}_j(t) = m_j^G(t)$, for $j = 1, 2$.

Figure 10 shows the disease dynamics in the targeted global social planner setting. Compared to the disease dynamics under laissez-faire, the globally optimal mitigation policy results in a less severe epidemic, as cumulative infections are reduced in both regions. Additionally, not only is aggregate social welfare higher (an unsurprising result

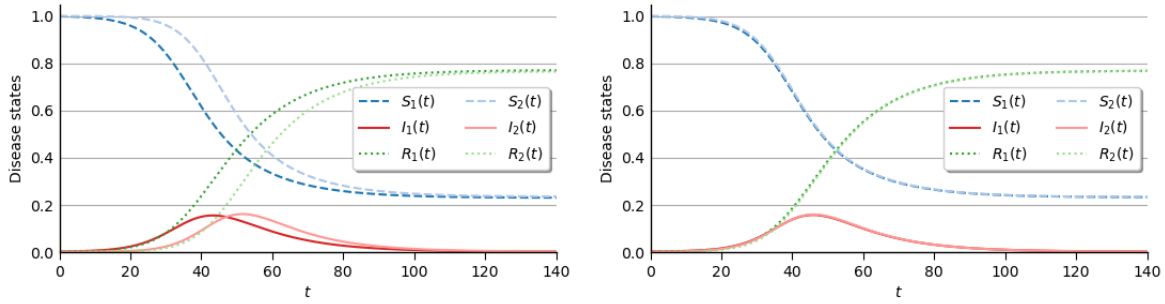

(a) Low ε. (b) High ε.

Figure 10: Disease dynamics in the targeted global social planner setting.

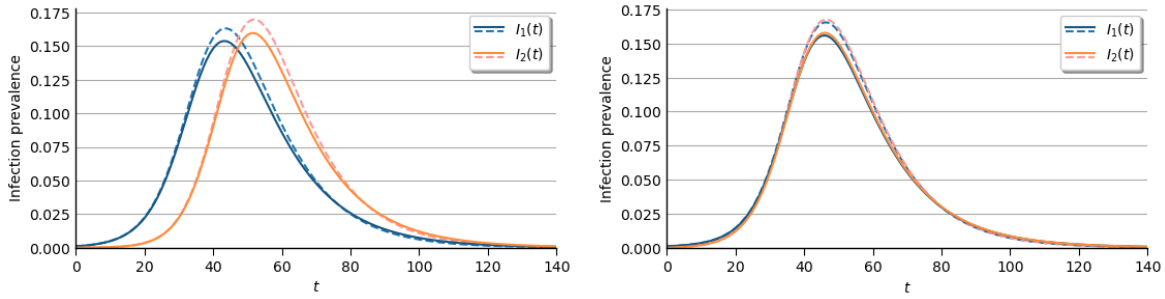

(a) Low ε. (b) High ε.

Figure 11: Infection prevalence in the targeted global social planner setting (solid lines) and the laissez-faire setting (dashed lines).

given the social planner's objective), but social welfare within each region is higher as well.

The improvements in cumulative infections and social welfare are approximately the same across the low and high ε specifications. This is due to the fact that the basic reproductive ratio is kept constant. Although mitigation in the stronger connected specification has a larger inter-regional externality, this is compensated by a reduction in its intra-regional externality. Thus, the total externalities internalised by the global social planner are roughly of the same magnitude. This is also evidenced in the weakly and strongly connected specifications showing similar impacts to infection prevalence and mitigation levels, as seen in Figures 11 and 12, respectively. The internalisation of the positive externalities from mitigation induces the global social planner to enact higher mitigation levels than under laissez-faire. For the parameter values used in these simulations, this is true at all points during the epidemic (i.e., $m_j^G(t) \geq m_j^*(t)$, $\forall j, t$). In other words, when non-cooperative individuals are allowed to independently choose how much they mitigate, they consistently exert less mitigation efforts than would be socially optimal, resulting in a more severe epidemic with more cumulative infections.

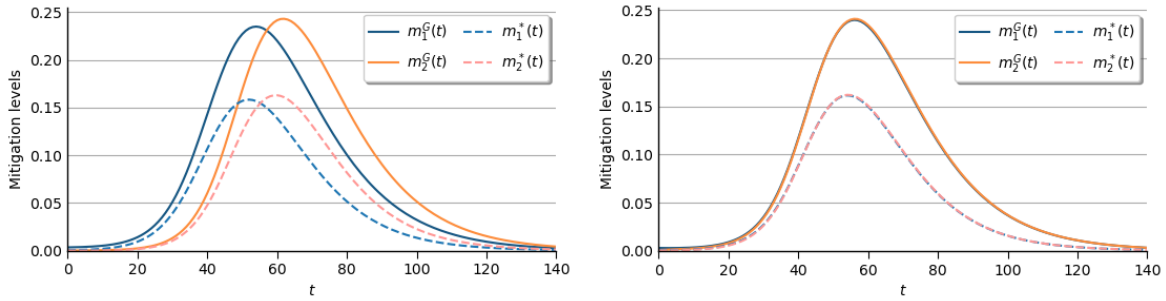

(a) Low ε. (b) High ε.

Figure 12: Mitigation levels in the targeted global social planner setting (solid lines) and the laissez-faire setting (dashed lines).

4.1.1 Uniform mitigation policies

In the above formulation of the global social planner's decision problem, we assume they are able to 'target' their mitigation policy, in that they can impose different mitigation levels to individuals in different regions at the same time. However, it may be the case that the global social planner is unable to impose different policies in neighbouring regions due to, for example, jurisdictional limitations or practical difficulties in implementation. In this scenario, which we will henceforth refer to as the 'uniform policy' setting, the global social planner would have the same objective function as before, but be limited by the constraint that $m_1(t) = m_2(t)$. Thus, denoting $m(t)$ as the mitigation level imposed in both regions at time $t \geq 0$, the *uniform* global social planner's decision problem is

$$\begin{aligned}\max_{m(t)\in[0,1]} \int_0^\infty e^{-\rho t} \Big[& S_1(t)\left(\pi_{\mathcal{S}} - \frac{c}{2}m(t)^2\right) + I_1(t)\pi_{\mathcal{I}} + R_1(t)\pi_{\mathcal{R}} \\ & + S_2(t)\left(\pi_{\mathcal{S}} - \frac{c}{2}m(t)^2\right) + I_2(t)\pi_{\mathcal{I}} + R_2(t)\pi_{\mathcal{R}}\Big] dt\end{aligned} \tag{4.7}$$

subject to, for $j = 1, 2$,

$$\dot{S}_j(t) = -\left(1 - m(t)\right)\beta S_j(t)\left(I_j(t) + \varepsilon I_{-j}(t)\right) \tag{4.8a}$$

$$\dot{I}_j(t) = \left(1 - m(t)\right)\beta S_j(t)\left(I_j(t) + \varepsilon I_{-j}(t)\right) - \gamma I_j(t) \tag{4.8b}$$

$$\dot{R}_j(t) = \gamma I_j(t) \tag{4.8c}$$

Similar to before, the uniform global social planner's current-value Hamiltonian is

$$\begin{aligned} H^U = {} & S_1(t)\left(\pi_{\mathcal{S}} - \frac{c}{2}m(t)^2\right) + I_1(t)\pi_{\mathcal{I}} + R_1(t)\pi_{\mathcal{R}} \\ & + S_2(t)\left(\pi_{\mathcal{S}} - \frac{c}{2}m(t)^2\right) + I_2(t)\pi_{\mathcal{I}} + R_2(t)\pi_{\mathcal{R}} \\ & - \lambda_{S_1}(t)\left(1-m(t)\right)\beta S_1(t)\left(I_1(t) + \varepsilon I_2(t)\right) \\ & - \lambda_{S_2}(t)\left(1-m(t)\right)\beta S_2(t)\left(I_2(t) + \varepsilon I_1(t)\right) \\ & + \lambda_{I_1}(t)\left[\left(1-m(t)\right)\beta S_1(t)\left(I_1(t) + \varepsilon I_2(t)\right) - \gamma I_1(t)\right] \\ & + \lambda_{I_2}(t)\left[\left(1-m(t)\right)\beta S_2(t)\left(I_2(t) + \varepsilon I_1(t)\right) - \gamma I_2(t)\right] \\ & + \lambda_{R_1}(t)\gamma I_1(t) \\ & + \lambda_{R_2}(t)\gamma I_2(t) \end{aligned} \tag{4.9}$$

for which the costate variables' laws of motion are, for $j = 1, 2$,

$$\begin{aligned} \dot{\lambda}_{S_j}(t) = {} & \lambda_{S_j}(t)\rho + \left(1-m(t)\right)\beta\left(I_j(t) + \varepsilon I_{-j}(t)\right)\left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right) \\ & - \pi_{\mathcal{S}} + \frac{c}{2}m(t)^2 \end{aligned} \tag{4.10a}$$

$$\begin{aligned} \dot{\lambda}_{I_j}(t) = {} & \lambda_{I_j}(t)\left(\rho + \gamma\right) + \left(1-m(t)\right)\beta S_j(t)\left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right) \\ & + \left(1-m(t)\right)\varepsilon\beta S_{-j}(t)\left(\lambda_{S_{-j}}(t) - \lambda_{I_{-j}}(t)\right) - \lambda_{R_j}(t)\gamma - \pi_{\mathcal{I}} \end{aligned} \tag{4.10b}$$

$$\dot{\lambda}_{R_j}(t) = \lambda_{R_j}(t)\rho - \pi_{\mathcal{R}} \tag{4.10c}$$

Taking the first order condition of the current-value Hamiltonian (4.9) and solving for $m(t)$ yields

$$m(t) = \frac{\beta S_1(t)\left(I_1(t) + \varepsilon I_2(t)\right)\left(\lambda_{S_1}(t) - \lambda_{I_1}(t)\right) + \beta S_2(t)\left(I_2(t) + \varepsilon I_1(t)\right)\left(\lambda_{S_2}(t) - \lambda_{I_2}(t)\right)}{c\left(S_1(t) + S_2(t)\right)} \tag{4.11}$$

We denote the socially optimal mitigation levels in the uniform policy setting as $m^U(t)$, being the solution to the necessary optimality condition (4.11), subject to $m(t) \in [0, 1]$ and the costate variables' laws of motion (4.10). The resulting SIR disease dynamics are then given by the metapopulation SIR model with mitigation (2.4) where $\bar{m}_j(t) = m^U(t)$, for $j = 1, 2$, and are illustrated in Figure 13.

Figures 14 and 15 illustrate that the impact of restricting the global social planner to uniform policies depends on the magnitude of ε. When ε is high and the regions are strongly connected, each region's disease states develop almost in tandem. Consequently, the socially optimal targeted mitigation levels are nearly equal between the regions. Thus, the global social planner being restricted to uniform policies does not have a significant impact on the mitigation levels imposed in each region.

When ε is low and the regions are weakly connected, there is differentiation in the regions' infection and optimal targeted mitigation levels. Thus, when restricted to uniform

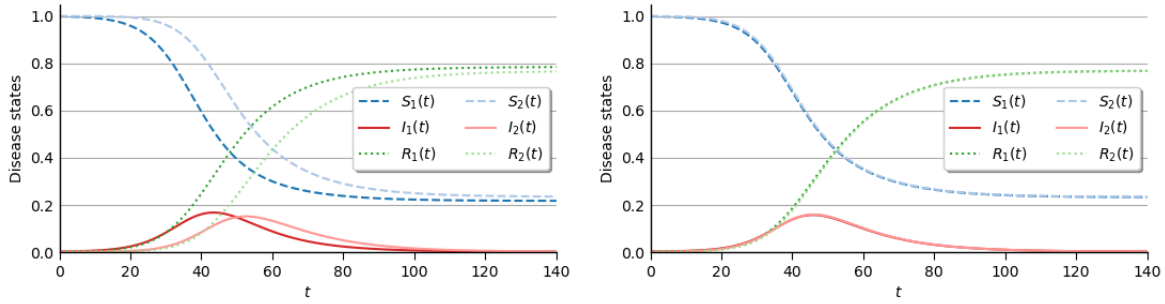

(a) Low ε. (b) High ε.

Figure 13: Disease dynamics in the uniform policy setting.

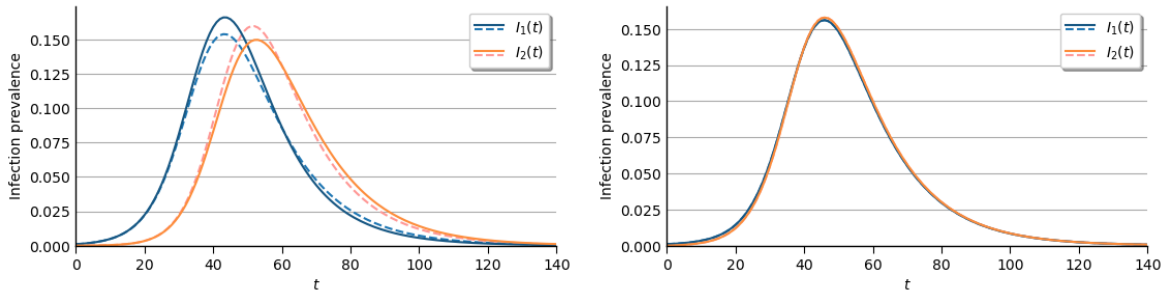

(a) Low ε. (b) High ε.

Figure 14: Infection prevalence under the global social planner's uniform policy (solid lines) and targeted policy (dashed lines).

policies, the global social planner sets mitigation levels by trading off between too much mitigation in one region and too little in another, resulting in $m^U(t)$ generally being between $m_1^G(t)$ and $m_2^G(t)$. There is a period, however, in which the uniform mitigation level is below the two targeted mitigation levels. This occurs around the time when the two targeted mitigation levels are equal ($m_1^G(t) = m_2^G(t)$). Correspondingly, the peak mitigation level under the uniform policy is noticeably lower than the peak mitigation level of either region under the targeted policy. These findings lead way to two insights. First, being able to tailor mitigation efforts to each region makes mitigation more effective as a tool. This induces its increased use by the global social planner under the targeted policy, explaining the lower peak mitigation level in the uniform policy setting. Second, the period in which the uniform mitigation level is less than both targeted mitigation levels highlights that, while a balancing between optimal targeted mitigation levels is involved, it is not the only component involved in the uniform global social planner's decision process. As the social planner implements the uniform mitigation level, the disease state paths differ from their paths in the optimal targeted setting. If the targeted global social planner were instead faced with these altered disease state paths, they would correspondingly update their optimal targeted mitigation levels. Indeed, as illustrated in Figure 16, when one continuously reevaluates the optimal targeted mitigation levels along

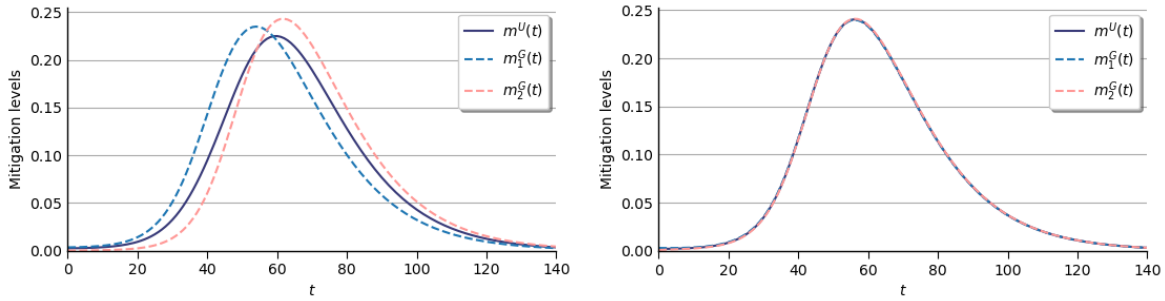

(a) Low ε. (b) High ε.

Figure 15: Mitigation levels under the global social planner's uniform policy (solid lines) and targeted policy (dashed lines).

the disease state paths of the uniform policy setting, $m^U(t)$ remains between the targeted mitigation levels.

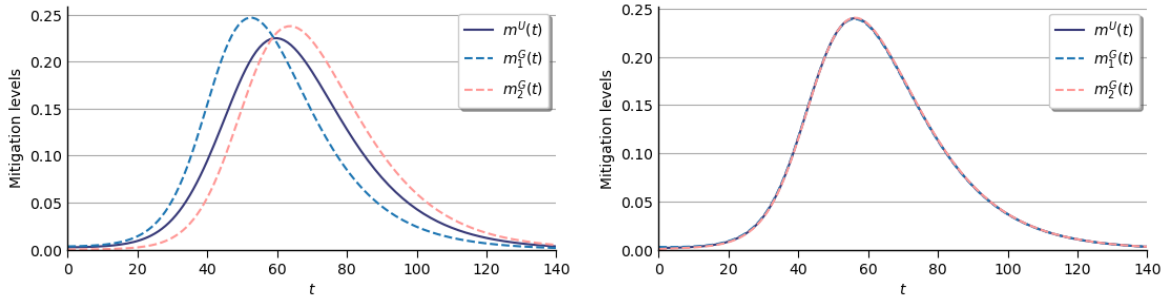

(a) Low ε. (b) High ε.

Figure 16: Mitigation levels under the uniform policy (solid lines) and the optimal targeted levels when evaluated along the uniform policy's disease state path (dashed lines).

The impacts of the uniform policy restriction on the disease dynamics and infection prevalence in each region is most easily seen in the case of weakly connected regions (Figure 14a). Relative to the targeted policy, mitigation levels in region 1 under the uniform policy are lower both leading up to and beyond the region's infection peak. As a result, infection prevalence in region 1 is 'spiked', in that it has a higher peak, but also dissipates quicker. The opposite is true for region 2. As mitigation levels are higher in region 2 during the epidemic's growth phase, the uniform policy results in a flattened infection prevalence curve with a longer duration epidemic. Overall, however, both regions have more cumulative infections under the uniform policy. Also, as is naturally expected given the uniform policy is a restriction on the global social planner's toolset, aggregate social welfare is lower than under the targeted policy.

As noted by Ndeffo Mbah and Gilligan (2011), a targeted disease control policy in which individuals are subjected to different rules depending on where they live may be viewed as socially inequitable. Thus, to counter such potential criticisms, the global social planner may opt to resort to a uniform mitigation policy. Doing so, however, results in

less equality between the regions in terms of epidemic severity and social welfare. In the uniform policy setting's low ε specification, the region with the initial small shock of infections experiences a noticeably higher peak level of infection. Correspondingly, the difference between the two regions in the cumulative number of infections is increased in this specification, with region 1 experiencing more infections over the course of the epidemic. This is similarly reflected in the social welfare, for which the disparity between the two regions is increased under the uniform policy for both ε specifications. The inequality in outcomes under the uniform policy is also greater than when mitigation is determined by laissez-faire. To summarise, although a uniform mitigation policy may be viewed as a more socially equitable tool, it can instead result in more unequal outcomes, both in terms of regional welfare and epidemic severity.

4.2 Delegated mitigation policies

We now turn to an alternate form of decision making, in which the decision to mitigate is not aggregated to the global level, but rather delegated to the regional level. In other words, mitigation levels are chosen in each region by a *regional* social planner, who does so in order to maximise the social welfare within their region.

Unlike in the previous cases of the global social planner or laissez-faire, this 'delegated policy' setting involves two uncoordinated social planners with differing objectives, for whom their mitigation decision has a non-negligible impact on the disease dynamics faced by the other social planner. Thus, the decision to mitigate by the regional social planners is a continuous-time differential game.

Due to the lack of analytical tractability already inherent in the SIR epidemiological model, we will focus on open-loop strategies. Open-loop strategies are Markovian (history-independent) strategies in which a social planner's choice of mitigation level depends solely on the current time t. By contrast, closed-loop strategies are Markovian strategies which depend on both the current time and disease states. As noted by Dockner et al. (2000), open-loop solutions are a subset of closed-loop solutions, and would involve each player announcing and committing to their mitigation policies at the start of the game. The commitment that lies within open-loop strategies can reflect the far-sightedness of social planners, who need to implement a mitigation policy to optimise welfare across an infinite planning horizon.

Denoting the mitigation level chosen by the regional social planner in region j at time t as $m_j(t)$, the decision problem for region j's social planner is written as

$$\max_{m_j(t)\in[0,1]} \int_0^\infty e^{-\rho t} \left[S_j(t) \left(\pi_{\mathcal{S}} - \frac{c}{2} m_j(t)^2 \right) + I_j(t)\pi_{\mathcal{I}} + R_j(t)\pi_{\mathcal{R}} \right] dt \tag{4.12}$$

subject to

$$\dot{S}_j(t) = -\left(1 - m_j(t)\right)\beta S_j(t)\left(I_j(t) + \varepsilon I_{-j}(t)\right) \tag{4.13a}$$
$$\dot{S}_{-j}(t) = -\left(1 - m_{-j}(t)\right)\beta S_{-j}(t)\left(I_{-j}(t) + \varepsilon I_j(t)\right) \tag{4.13b}$$
$$\dot{I}_j(t) = \left(1 - m_j(t)\right)\beta S_j(t)\left(I_j(t) + \varepsilon I_{-j}(t)\right) - \gamma I_j(t) \tag{4.13c}$$
$$\dot{I}_{-j}(t) = \left(1 - m_{-j}(t)\right)\beta S_{-j}(t)\left(I_{-j}(t) + \varepsilon I_j(t)\right) - \gamma I_{-j}(t) \tag{4.13d}$$
$$\dot{R}_j(t) = \gamma I_j(t) \tag{4.13e}$$
$$\dot{R}_{-j}(t) = \gamma I_{-j}(t) \tag{4.13f}$$

Denoting the costate variable of disease compartment τ_k for the social planner of region j as $\lambda^j_{\tau_k}(t)$, region j's social planner's current-value Hamiltonian is

$$\begin{aligned} H_j = {} & S_j(t)\left(\pi_{\mathcal{S}} - \frac{c}{2}m_j(t)^2\right) + I_j(t)\pi_{\mathcal{I}} + R_j(t)\pi_{\mathcal{R}} \\ & - \lambda^j_{S_j}(t)\left(1 - m_j(t)\right)\beta S_j(t)\left(I_j(t) + \varepsilon I_{-j}(t)\right) \\ & - \lambda^j_{S_{-j}}(t)\left(1 - m_{-j}(t)\right)\beta S_{-j}(t)\left(I_{-j}(t) + \varepsilon I_j(t)\right) \\ & + \lambda^j_{I_j}(t)\left[\left(1 - m_j(t)\right)\beta S_j(t)\left(I_j(t) + \varepsilon I_{-j}(t)\right) - \gamma I_j(t)\right] \\ & + \lambda^j_{I_{-j}}(t)\left[\left(1 - m_{_j}(t)\right)\beta S_{-j}(t)\left(I_{-j}(t) + \varepsilon I_j(t)\right) - \gamma I_{-j}(t)\right] \\ & + \lambda^j_{R_j}(t)\gamma I_j(t) \\ & + \lambda^j_{R_{-j}}(t)\gamma I_{-j}(t) \end{aligned} \tag{4.14}$$

The current-value costate variables' laws of motion are then given by

$$\begin{aligned} \dot{\lambda}^j_{S_j}(t) = {} & \lambda^j_{S_j}(t)\rho + \left(1 - m_j(t)\right)\beta\left(I_j(t) + \varepsilon I_{-j}(t)\right)\left(\lambda^j_{S_j}(t) - \lambda^j_{I_j}(t)\right) \\ & - \pi_{\mathcal{S}} + \frac{c}{2}m_j(t)^2 \end{aligned} \tag{4.15a}$$
$$\dot{\lambda}^j_{S_{-j}}(t) = \lambda^j_{S_{-j}}(t)\rho + \left(1 - m_{-j}(t)\right)\beta\left(I_{-j}(t) + \varepsilon I_j(t)\right)\left(\lambda^j_{S_{-j}}(t) - \lambda^j_{I_{-j}}(t)\right) \tag{4.15b}$$
$$\begin{aligned} \dot{\lambda}^j_{I_j}(t) = {} & \lambda^j_{I_j}(t)\left(\rho + \gamma\right) + \left(1 - m_j(t)\right)\beta S_j(t)\left(\lambda^j_{S_j}(t) - \lambda^j_{I_j}(t)\right) \\ & + \left(1 - m_{-j}(t)\right)\varepsilon\beta S_{-j}(t)\left(\lambda^j_{S_{-j}}(t) - \lambda^j_{I_{-j}}(t)\right) - \lambda^j_{R_j}(t)\gamma - \pi_{\mathcal{I}} \end{aligned} \tag{4.15c}$$
$$\begin{aligned} \dot{\lambda}^j_{I_{-j}}(t) = {} & \lambda^j_{I_{-j}}(t)\left(\rho + \gamma\right) + \left(1 - m_{-j}(t)\right)\beta S_{-j}(t)\left(\lambda^j_{S_{-j}}(t) - \lambda^j_{I_{-j}}(t)\right) \\ & + \left(1 - m_j(t)\right)\varepsilon\beta S_j(t)\left(\lambda^j_{S_j}(t) - \lambda^j_{I_j}(t)\right) - \lambda^j_{R_{-j}}(t)\gamma \end{aligned} \tag{4.15d}$$
$$\dot{\lambda}^j_{R_j}(t) = \lambda^j_{R_j}(t)\rho - \pi_{\mathcal{R}} \tag{4.15e}$$
$$\dot{\lambda}^j_{R_{-j}}(t) = \lambda^j_{R_{-j}}(t)\rho \tag{4.15f}$$

As only open-loop solutions are considered, the social planner for region j takes $m_{-j}(t)$ as given during their optimisation, and selects $m_j(t)$ so as to maximise their current-value Hamiltonian (4.14). Similar to before, differentiating (4.14) with respect to $m_j(t)$ and

assuming $S_j(t) > 0$ gives the necessary optimality condition

$$m_j(t) = \frac{\beta\left(I_j(t) + I_{-j}(t)\right)\left(\lambda^j_{S_j}(t) - \lambda^j_{I_j}(t)\right)}{c} \tag{4.16}$$

Before continuing with the discussion of the outcomes in the delegated policy setting, we will take a moment to compare the decision problem and corresponding solutions with that of the targeted global social planner setting. The costate variables's laws of motion (4.15) and the optimality condition (4.16) in the delegated policy setting appear almost identical to those of the targeted global social planner (4.5), (4.6). However, as the utility flows earned in the opposing region are not considered by region j's social planner, their costate variables for the opposing region move differently to account for this. Thus, for each disease state τ, the rate of change of the inter-regional costate variable $\lambda^j_{\tau_{-j}}(t)$ contains the same components as its intra-regional counterpart, $\lambda^j_{\tau_j}(t)$, except for the utility flow earned by individuals in $\tau_{-j}(t)$. Similarly, as discussed further in Appendix A.1, the transversality conditions for the inter-regional costate variables $\lambda^j_{\tau_{-j}}(t)$ do not value the continuation payoffs earned by the individuals in $\tau_{-j}(t)$, unlike their intra-regional counterparts.

We denote the (regionally) optimal open-loop mitigation levels in region j in the delegated policy setting as $m^R_j(t)$, being the solution to the necessary optimality condition (4.16), subject to $m_j(t) \in [0, 1]$ and the costate variables' laws of motion (4.15). The resulting SIR disease dynamics are then given by the metapopulation SIR model with mitigation (2.4) where $\bar{m}_j(t) = m^R_j(t)$, for $j = 1, 2$. Figure 17 illustrates the disease dynamics in the delegated policy setting, whereas Figure 18 shows the infection prevalence levels compared with those in the targeted global social planner setting.

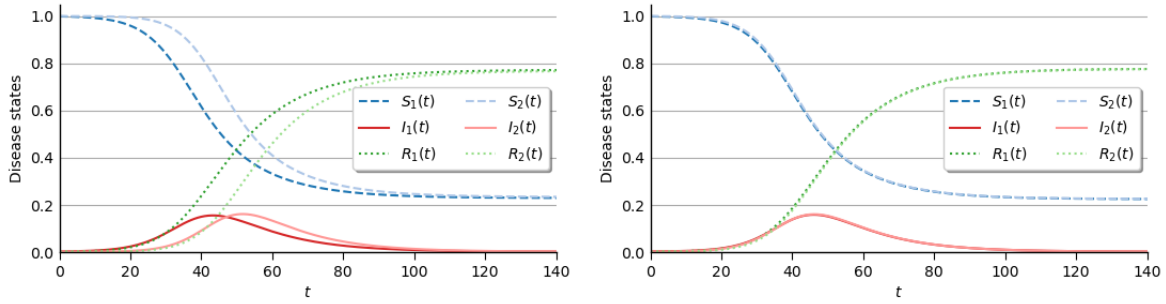

(a) Low ε. (b) High ε.

Figure 17: Disease dynamics in the delegated policy setting.

As shown in Figure 19, the regional social planner generally imposes lower mitigation levels than the global planner. This result makes intuitive sense, as the regional social planner only internalises mitigation's intra-regional externalities, whereas the global social planner also internalises the inter-regional externalities.

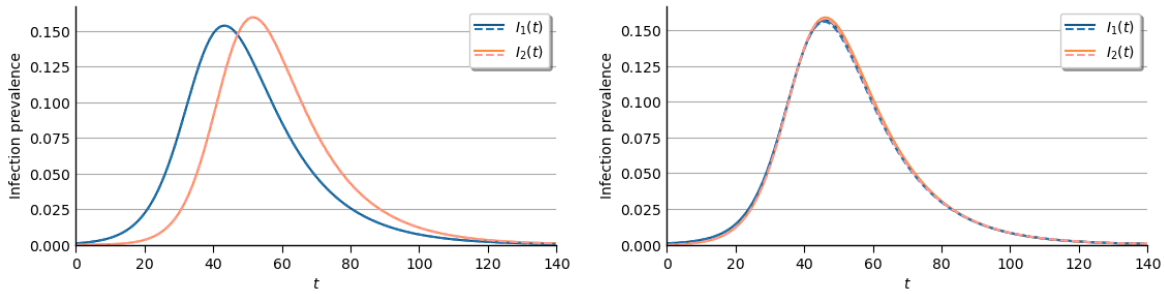

(a) Low ε. (b) High ε.

Figure 18: Infection prevalence in the delegated policy setting (solid lines) and the targeted global social planner setting (dashed lines).

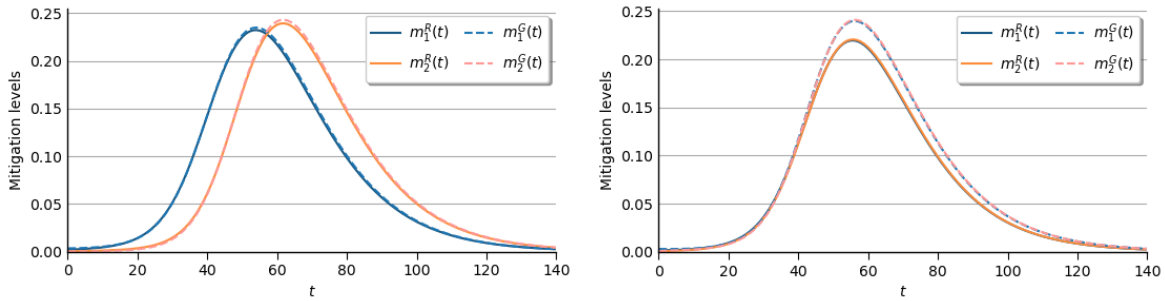

(a) Low ε. (b) High ε.

Figure 19: Mitigation levels in the delegated policy setting (solid lines) and the targeted global social planner setting (dashed lines).

Although the regional social planner's mitigation levels are generally lower than those of the global social planner, there is a period during the epidemic where this is not the case. Starting near the beginning of the epidemic, the regional social planner imposes stricter mitigation measures than the global social planner. This continues until mitigation levels increase at their most rapid rate, after which the regional social planner enacts less mitigation, and continues to do so into the infinite horizon. One possible explanation for this comes from the inter-regional strategic nature of mitigation, discussed in Section 3.1. As noted then, increased levels of mitigation in the neighbouring region induce individuals in the home region to mitigate less in the earlier stages of the epidemic, and more in the end stages. Thus, as we transition from the targeted global social planner setting to the delegated policy setting, the opposite is true. Region 2's social planner lessens mitigation (due to no longer internalising the inter-regional externalities), inducing region 1's social planner to increase mitigation at the beginning of the epidemic, and decrease mitigation towards the end. Of course, this is reciprocally true for region 1's impact on region 2. Thus, relative to the targeted global social planner setting, the inter-regional strategic effect of mitigation causes an upwards pressure on a regional social planner's mitigation efforts at the beginning of the epidemic, resulting in higher mitigation levels.

Overall, for both specifications of ε, the delegated policy setting results in less social

welfare and more cumulative infections in both regions. The magnitude of this impact, however, changes with ε. For strongly connected regions, the transition from the targeted global social planner setting to the delegated policy setting has a larger impact in all respects (mitigation levels, infections, and social welfare). This dependency is explained by the fact that the inter-regional externalities and strategic effects are both greater when ε is large.

4.3 Decision making delegation under uniform policies

Examples as to why a global social planner may be restricted to only implementing uniform mitigation policies across their jurisdiction are discussed in Section 4.1.1. However, it may be the case that the uniform policy-restricted global social planner can impose region-specific mitigation levels by instead delegating policy determination to the underlying regional social planners. These regional social planners would then impose mitigation levels specific to their respective jurisdictions. The question then arises as to when a social planner would find such a delegation to be beneficial. To answer this, we compare the outcomes of the uniform policy setting with those of the delegated policy setting, which is presented in Figures 20 and 21.

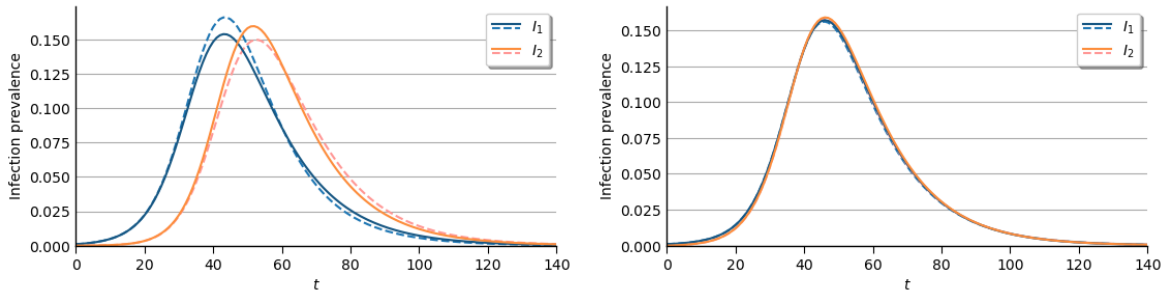

(a) Low ε.

(b) High ε.

Figure 20: Infection prevalence in the delegated policy setting (solid lines) and the uniform policy setting (dashed lines).

In the case of low ε, the comparison between the delegated policy and uniform policy settings (Figures 20a, 21a) looks similar to the comparison between the targeted global social planner and uniform policy settings (Figures 14a, 15a). This is because, due to the regions' weak coupling, the magnitude of mitigation's inter-regional externalities and strategic effects are low. Thus, the mitigation levels and disease state paths under the delegated policy are similar to those under the targeted global social planner. Furthermore, as there is regional heterogeneity in the disease state dynamics caused by the low value of ε, there is a high cost of resorting to uniform policies. As a result, cumulative infections are lower and total social welfare is higher in the delegated policy setting than in the uniform policy setting. Thus, if the regions under their domain are weakly connected, a

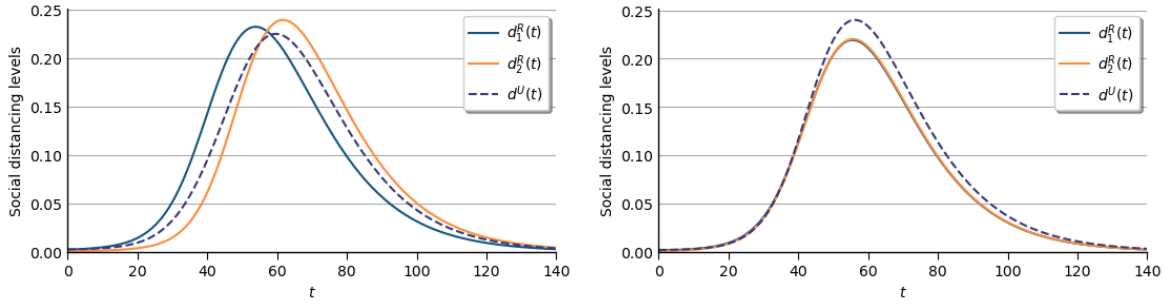

(a) Low ε. (b) High ε.

Figure 21: Mitigation levels in the delegated policy setting (solid lines) and the uniform policy setting (dashed lines).

global social planner restricted to uniform policies would likely benefit from delegating the mitigation decision making to the underling regional social planners.

For the high ε specification, the comparison of the delegated policy and uniform policy settings (Figures 20b, 21b) instead looks like the comparison between the delegated policy and targeted global social planner settings (Figures 18b, 19b). As the regions in this specification are strongly connected, the inter-regional externalities are also high. This means that the social welfare cost of delegating to uncoordinated regional social planners is large. On the other hand, the regional homogeneity in disease state dynamics means the cost of uniform policies is low. This is seen in the similarity in disease dynamics and mitigation levels between the targeted and uniform policies of the global social planner. Therefore, compared to the uniform policy setting, there are more cumulative infections and less social welfare under the delegated policy setting. Contrary to the specification of low ε, when a global social planner's underlying regions are strongly connected, they would benefit from retaining the decision-making power and implementing a uniform mitigation policy.

Letting a global social planner set targeted policies is always the first-best setting; but whether a particular value of ε entails that a delegated policy is preferred to a uniform policy depends on the values of the model's other parameters. However, as illustrated in Figure 22, there will be a threshold value of ε for which the welfare under the two inefficient policy settings are equal. For values of ε less than this threshold, the delegated policy results in higher welfare—approaching that of the first-best targeted global social planner—whereas the uniform policy has a welfare cost that increases as $\varepsilon \to 0$. On the other hand, as $\varepsilon \to 1$, the uniform policy results in welfare close to the first-best targeted global social planner, while the delegated policy results in larger welfare costs. Similarly, there is a threshold value of ε for which the cumulative infections are the same under the two inefficient policy settings. Lower values of ε lead to the delegated policy resulting in fewer infections, whereas higher values of ε lead to the uniform policy having fewer infections. Again, when $\varepsilon \to 0$, the delegated policy approximates the first-best outcome

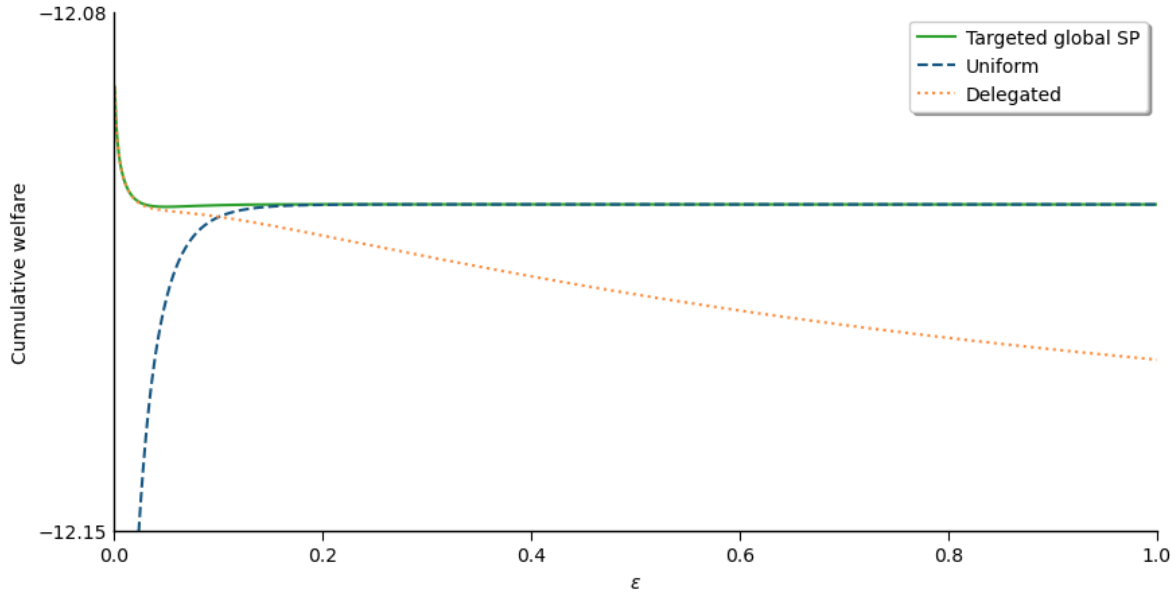

(a) Cumulative welfare.

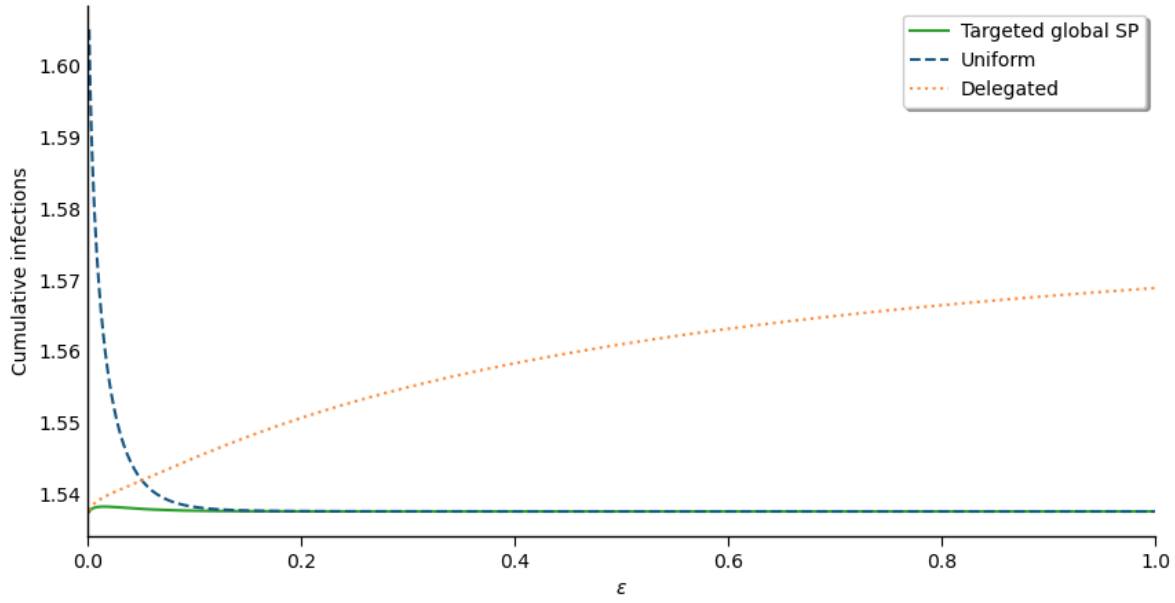

(b) Cumulative infections.

Figure 22: Total cumulative welfare (a) and cumulative infections (b) for different values of ε in the targeted global social planner (solid lines), uniform policy (dashed lines), and delegated policy (dotted lines) settings.

while the uniform policy results in high societal costs. When $\varepsilon \to 1$, the uniform policy approximates the first-best outcome while the delegated policy has high societal costs.

5 Travel restrictions

So far, we have assumed that the individuals and social planners of this model are unable to change the level of inter-regional interaction and, thus, take ε as given. What if instead, travel restrictions could be imposed, thereby reducing ε and the degree of interaction between regions? In this section, we calculate the decision maker's willingness to pay (WTP) for a marginal decrease in ε from a given level $\hat{\varepsilon} \in (0, 1]$. To do this, we compare the regional and total welfare from a simulation where $\varepsilon = \hat{\varepsilon}$ with that from a simulation where $\varepsilon = \hat{\varepsilon} - 0.001$. The welfare in region j for a certain value of ε is calculated as

$$W_j(\varepsilon) = \int_0^\infty e^{-\rho t} \left[S_j(t) \left(\pi_{\mathcal{S}} - \frac{c}{2} \bar{m}_j(t)^2 \right) + I_j(t)\pi_{\mathcal{I}} + R_j(t)\pi_{\mathcal{R}} \right] dt \tag{5.1}$$

and the total welfare is $W(\varepsilon) = W_1(\varepsilon) + W_2(\varepsilon)$. The WTP in region j to decrease ε by 0.001 from $\hat{\varepsilon}$ is $WTP_j(\hat{\varepsilon}) = W_j(\hat{\varepsilon} - 0.001) - W_j(\hat{\varepsilon})$. The total WTP is then the sum of the two regional WTPs.

As before, we adjust β for all simulations in this section in order to maintain a constant basic reproductive ratio R_0, irrespective of ε. Thus, a simulation using a value of ε will have a transmission rate $\beta(\varepsilon) = \frac{R_0\gamma}{1+\varepsilon}$. This is done so that the severity of the disease is not altered. In the case of travel restrictions, this entails that when ε is reduced, individuals respond by replacing their 'lost' inter-regional interactions with additional intra-regional interactions. Thus, the total number of interactions individuals maintain stays the same, as does the disease's basic reproductive ratio R_0. In other words, we keep constant an individual's total number of contacts, such that a change in inter-regional connectedness only influences the distribution of their contacts across the home and neighbouring region. If we did not make this adjustment, the calculated WTP to reduce ε would be biased with the WTP to reduce the disease's severity.

5.1 WTP without mitigation

As a benchmark, the regional and total WTP in the case of the purely mechanistic model with no mitigation is shown in Figure 23. Since there is no mitigation or decision making in this setting, the change in welfare caused by a decrease in ε is entirely driven by the resulting changes in the disease dynamics. As the severity of the disease remains constant, only the epidemic timings are changed and, thus, entirely explain the WTP curves.[12] Namely, when ε is decreased and the regions become less interconnected, it takes longer for the disease to meaningfully spread to region 2. As the epidemic and its associated costs are delayed for the second region, this is beneficial due to discounting. Hence, the WTP in region 2 is positive. Conversely, a lower ε means that the initial infection in region 1 is effectively less 'diluted' and region 1 undergoes the epidemic sooner. Again,

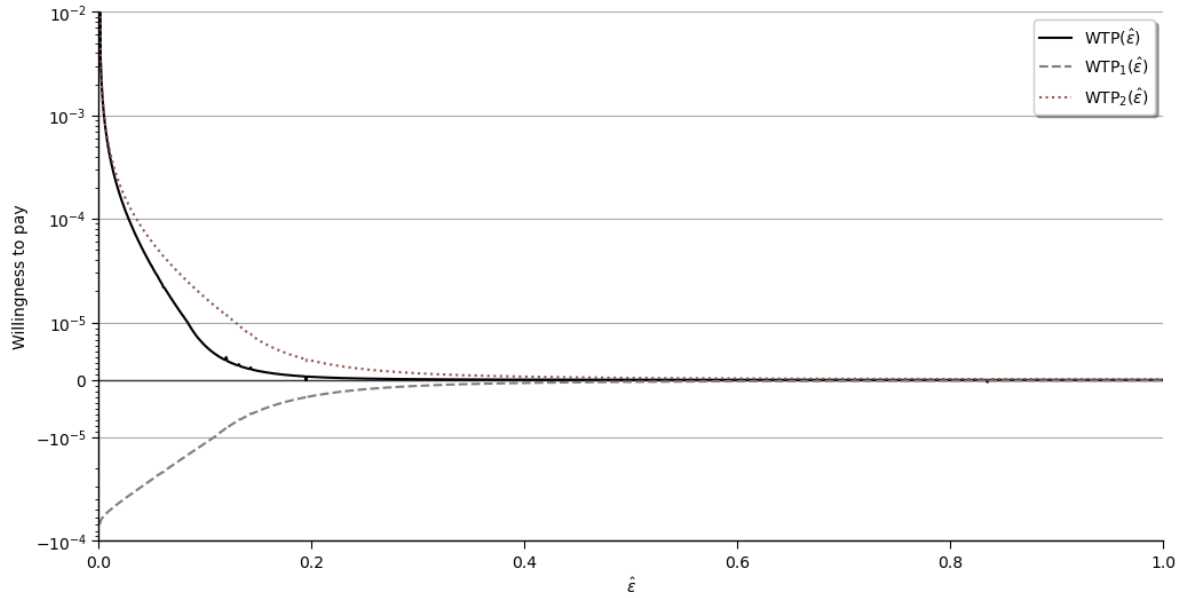

Figure 23: Regional and total WTP to decrease ε by 0.001 in the mechanistic model with no mitigation.

due to discounting, this makes region 1 worse off and gives them a negative regional WTP. In total, as the changes in the epidemic timing for region 1 are marginal compared to the epidemic delay in region 2, the total WTP is positive for all $\hat{\varepsilon}$. Furthermore, the regional and total WTP are all monotone and experience their largest magnitudes as $\hat{\varepsilon} \to 0$. This is because for low $\hat{\varepsilon}$, the difference in epidemic timings between the two regions is more pronounced and an absolute change in ε of 0.001 is relatively larger. Hence, the marginal decrease in ε has a larger effect. This is contrasted with larger values of $\hat{\varepsilon}$, in which the initial infection almost immediately spreads to region 2. Then, a marginal decrease in ε has a negligible impact on disease dynamics, and a correspondingly negligible impact on welfare.

5.2 WTP with endogenous mitigation

When considering the WTP in the model with endogenous mitigation in the laissez-faire setting (Figure 24), the WTP curves exhibit similar qualitative traits compared to the mechanistic benchmark model. This is because the epidemic timing effects from before are still present. However, there are notable differences, particularly in the non-monotonicity of the regional WTP. As before, when $\hat{\varepsilon}$ is large, the WTP in region 1 is negative and is increasing towards 0 as $\hat{\varepsilon} \to 1$. However, as $\hat{\varepsilon}$ becomes low and we move to the left in the graph, the slope of region 1's WTP curve changes sign and starts increasing with smaller values of $\hat{\varepsilon}$. This continues until the region's WTP eventually becomes positive. The reason for this can be found by applying the dynamic envelope theorem to the

[12]This can be seen by comparing the WTP for decreasing discount rates ρ (see Appendix D.1).

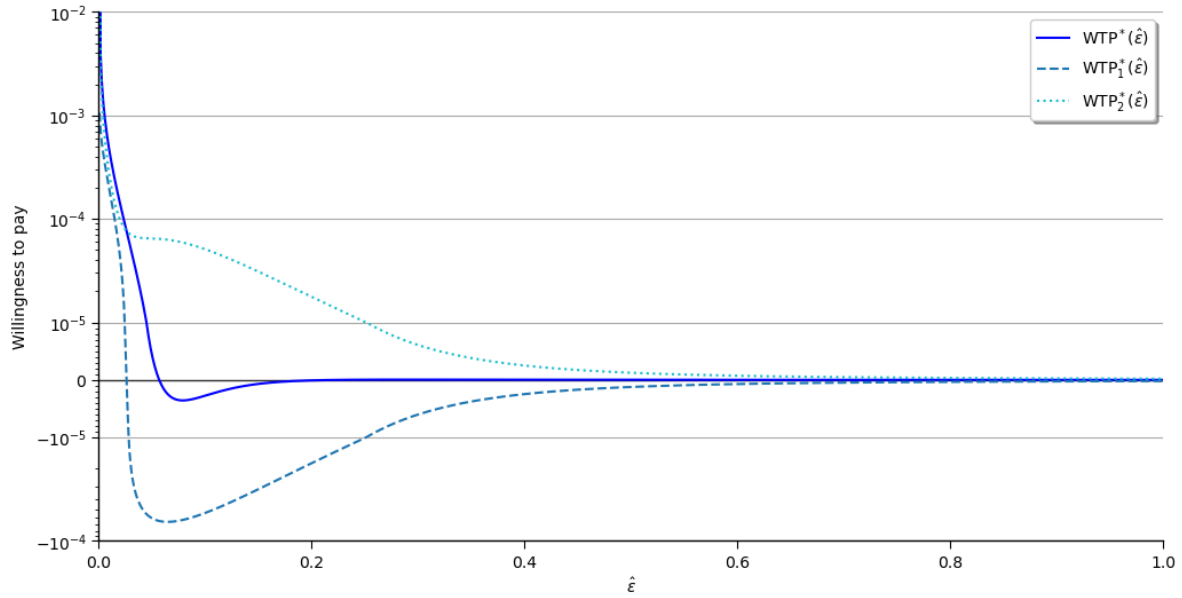

Figure 24: Regional and total WTP to decrease ε by 0.001 with laissez-faire mitigation.

individual's current-value Hamiltonian (3.2).[13] By doing so, we can calculate the WTP of an individual i in region j to marginally decrease ε as

$$\begin{aligned} WTP_j^* = \int_0^\infty e^{-\rho t} \Bigg[& \left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right) \\ & \times \left(1 - m_j^*(t)\right) \beta p^i_{S_j}(t) \\ & \times \left(\frac{I_{-j}(t) - I_j(t)}{1+\varepsilon} + \frac{\partial I_j}{\partial \varepsilon}(t) + \varepsilon \frac{\partial I_{-j}}{\partial \varepsilon}(t) \right) \Bigg] dt \end{aligned} \tag{5.2}$$

Note that this formulation (5.2) is an integral of the product of three time-varying factors. The first factor

$$\left(\lambda_{S_j}(t) - \lambda_{I_j}(t)\right) \tag{5.3}$$

captures the value of avoiding infection at time $t \geq 0$. As noted in Section 3, this expression is positive and increasing throughout the epidemic.

The second factor

$$\left(1 - m_j^*(t)\right) \beta p^i_{S_j}(t) \tag{5.4}$$

captures the rate of new infections per unit of effective infection prevalence, which is the term $(I_j(t) + \varepsilon I_{-j}(t))$. As is the case for the first factor, this factor is also positive $\forall t \in [0, \infty)$.

[13] See Caputo (2005, pp. 232–235) for a derivation of the dynamic envelope theorem in continuous-time optimal control problems. Note that our result slightly differs from the one suggested by Caputo, in that it includes the terms $\frac{\partial I_j(t)}{\partial \varepsilon}$ and $\varepsilon \frac{\partial I_{-j}(t)}{\partial \varepsilon}$. This is due to the fact that the equilibrium disease state paths change in response to a reduction in ε, but the individual takes them as given and does not maximise their current-value Hamiltonian with respect to them. Thus, $\frac{\partial H_j^i}{\partial I_j(t)}, \frac{\partial H_j^i}{\partial I_{-j}(t)} \neq 0$.

The last factor, which we term the *infection differential effect*, is

$$\left(\frac{I_{-j}(t) - I_j(t)}{1+\varepsilon} + \frac{\partial I_j}{\partial \varepsilon}(t) + \varepsilon \frac{\partial I_{-j}}{\partial \varepsilon}(t)\right) \tag{5.5}$$

This factor expresses the reduction in effective infection prevalence at time t caused by a reduction in ε and, unlike the first two factors, can be either positive or negative. The infection differential effect captures the idea that, when evaluated along the equilibrium disease state paths, a reduction in ε benefits an individual by reducing their interaction with and susceptibility to infected individuals in the neighbouring region, but simultaneously increases their interaction with and susceptibility to infected individuals within their own region. In general, if the neighbouring region has more infections than the home region, then the infection differential effect is positive, and vice versa.[14] In the early stages of the epidemic, when $I_1(t) > I_2(t)$, the infection differential effect is negative (positive) for region 1 (2). This reverses later on when infections grow in region 2 and $I_2(t) > I_1(t)$. Furthermore, for large values of $\hat{\varepsilon}$, infection levels in the two regions move almost identically, so the infection differential effect is approximately zero. When $\hat{\varepsilon}$ is low, the difference between $I_1(t)$ and $I_2(t)$ can be large.

The infection differential effect is the main source of the different characteristics of the WTP with endogenous mitigation. Relative to the mechanistic model, the inclusion of mitigation alters the disease dynamics: epidemic durations are extended, but regional infections still reach their peak prevalence at approximately the same time. Thus, compared to the WTP in the mechanistic model, the period of positive infection differential effects for region 1 starts at the same time, but lasts longer. When $\hat{\varepsilon}$ is low, these extended benefits are substantial enough to overcome the discounting effect and provide a net positive WTP in region 1. This is demonstrated in Figure 25. By the end of the epidemic, the cumulative discounted infection differential effect is positive in the laissez-faire setting, whereas it is approximately 0 in the mechanistic model without mitigation.

In other words, a reduction in ε makes those in the initially infected region 1 worse off at first, as they interact more with themselves and less with healthy individuals in region 2. As the disease spreads to region 2 and infections decline in region 1, the travel restrictions now provide a benefit, as susceptibles in region 1 now interact more with the recovered individuals in their own region, rather than the infected individuals in the neighbouring region. However, as people's decision to mitigate extends the epidemic's duration, this period of benefits lasts for a longer period of time. Thus, when regions are weakly connected, even individuals in an initially infected region may benefit from travel restrictions with initially uninfected regions.

[14] It should be noted that the infection differential effect's additional terms $\frac{\partial I_j}{\partial \varepsilon}(t)$ and $\varepsilon \frac{\partial I_{-j}}{\partial \varepsilon}(t)$ are generally marginal in magnitude compared to the primary term $\frac{I_{-j}(t) - I_j(t)}{1+\varepsilon}$. Thus, it is mainly this primary term that determines the sign and magnitude of the infection differential effect.

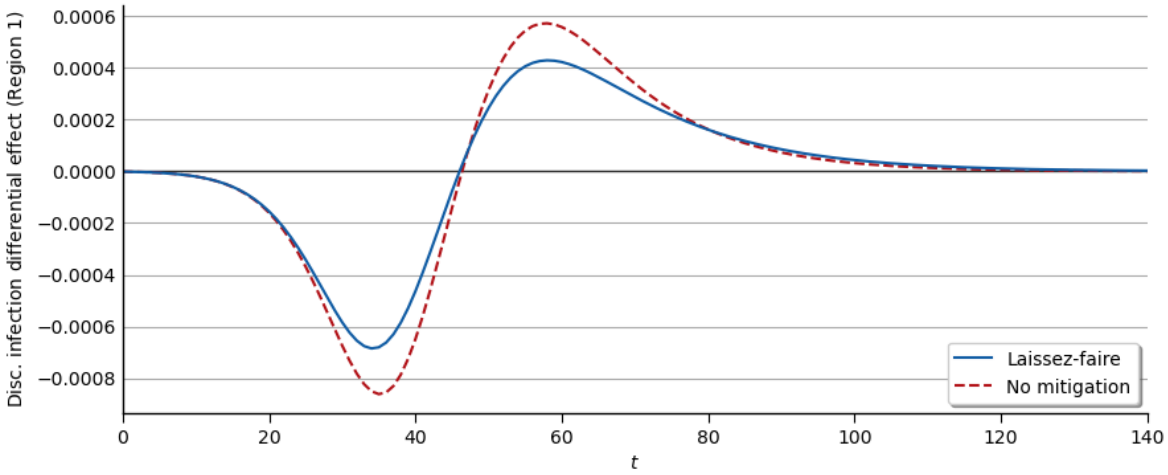

(a) Discounted infection differential effect.

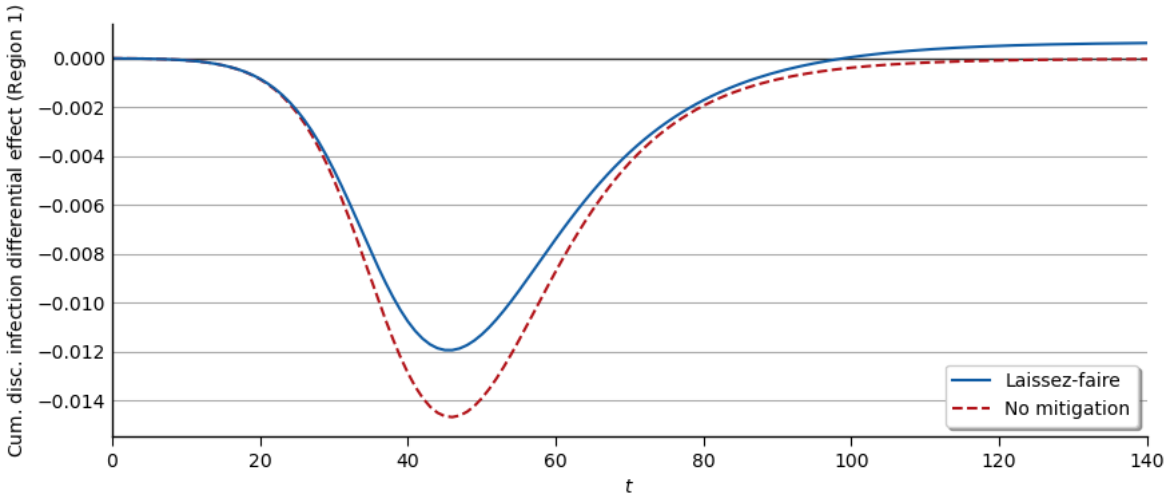

(b) Cumulative discounted infection differential effect.

Figure 25: (Cumulative) discounted infection differential effect for Region 1 under laissez-faire mitigation (solid blue line) and the mechanistic model without mitigation (dashed red line), when $\hat{\varepsilon} = 0.005$.

The opposite effect occurs for region 2: reductions in ε under laissez-faire mitigation leads to extended periods of costs when $I_2(t) > I_1(t)$. This effect puts a downward pressure for certain low values of $\hat{\varepsilon}$, causing the $WTP_2^*(\hat{\varepsilon})$ curve to become almost flat. This is even more apparent when using a lower discount rate ρ, as it is easier for the infection differential effect to overcome the discounting effect. Appendix D.2 shows that for lower values of ρ, it is possible for $WTP_2^*(\hat{\varepsilon})$ to become negative.

The combination of the infection differential effects in the two regions cause the total $WTP^*(\hat{\varepsilon})$ to be approximately 0 for large values of $\hat{\varepsilon}$, negative for a range of smaller $\hat{\varepsilon}$, and positive for very low values of $\hat{\varepsilon}$. This differs from the predictions under the mechanistic model, particularly in that the total WTP can be negative. Additionally, with lower values of ρ, the non-monotonicity of the total $WTP^*(\hat{\varepsilon})$ curve is more pronounced (see Appendix

D.2), leading to a larger range and magnitude of negative $WTP^*(\hat{\varepsilon})$. This suggests that if mitigation is determined endogenously by individuals, certain travel restrictions may actually be detrimental to society as a whole. Furthermore, this can be more likely to happen if low discount rates are used.

Similar results are found when observing the total WTP in the targeted social planner setting, as shown in Figure 26. This is because the same effects from the laissez-faire setting are present. Thus, when $\hat{\varepsilon}$ is large, $WTP^G(\hat{\varepsilon})$ is approximately 0. For lower values of $\hat{\varepsilon}$, $WTP^G(\hat{\varepsilon}) < 0$. As we move further to the left of the graph and $\hat{\varepsilon} \to 0$, $WTP^G(\hat{\varepsilon})$ becomes positive. These results suggest that if a global social planner is setting mitigation policies, they may be made worse off by imposing travel restrictions if they are not severe enough.

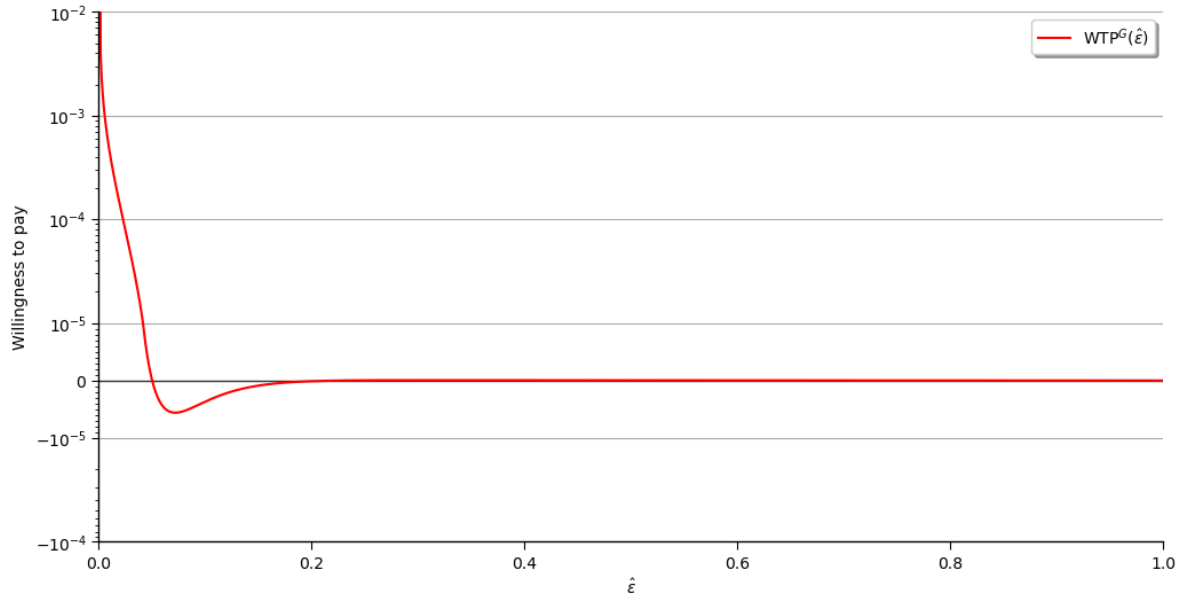

Figure 26: Total WTP to decrease ε by 0.001 in the targeted global social planner setting.

In the delegated policy setting, Figure 27 shows that the regional and total WTP curves have similar shapes compared to the laissez-faire and targeted global social planner settings, but are generally shifted upwards. This is because there is an additional positive effect in this setting: when ε decreases, not only do the disease dynamics change, but so does the strategic setting between the regional social planners. As their regions become less connected, they have less of an impact on each other, and the magnitude of the inter-regional externalities decreases. This induces them to behave more closely to the global social planner, resulting in better outcomes for both regions. Thus, there is a positive bias in the WTP. With the exception of $WTP_1^R(\hat{\varepsilon})$ for a range of small $\hat{\varepsilon}$, the regional and total WTP is positive, even for large values of $\hat{\varepsilon}$.

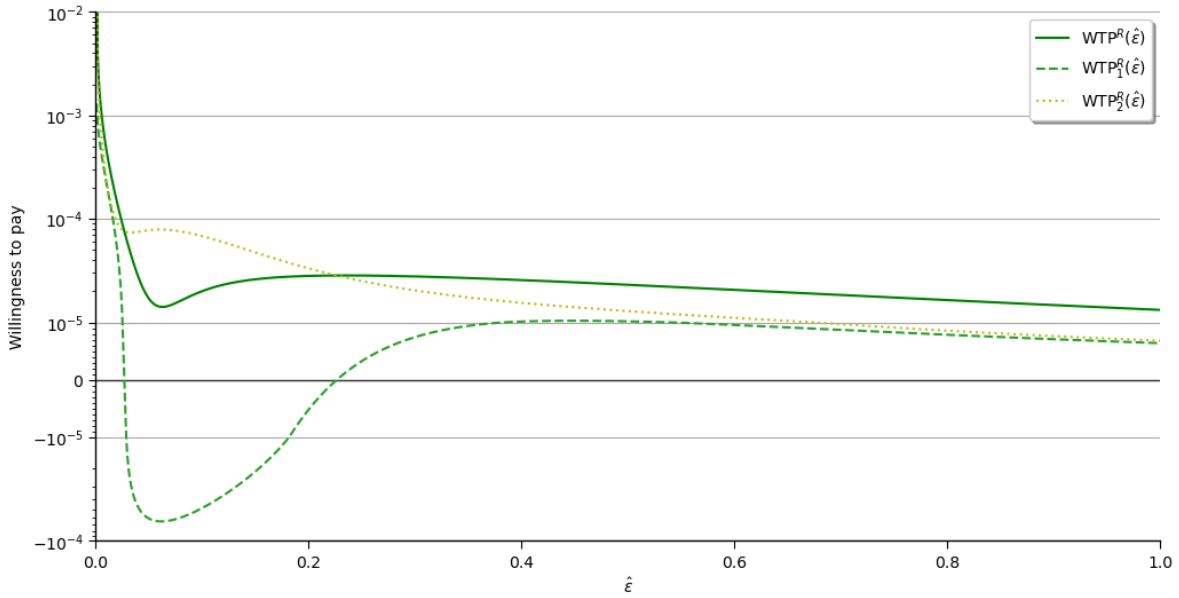

Figure 27: Regional and total WTP to decrease ε by 0.001 in the delegated policy setting.

6 Conclusion

In this paper, we use a metapopulation model to analyse the determination and impact of mitigation on the spatio-temporal dynamics of an infectious disease. We consider both the outcome of laissez-faire mitigation (which leads to uninternalised intra-regional externalities) and delegated mitigation (which leads to uninternalised inter-regional externalities). We emphasise the role of policy coordination and policy targeting, respectively, in achieving socially optimal outcomes. While lack of coordination or lack of targeting both induce inefficiencies relative to the global social optimum, the magnitudes of these depend on the strength of coupling between regions.

Our analysis can be enriched in several different directions. First, our analysis has focused on the case where only susceptible individuals can choose to mitigate, for the sake of comparability with the laissez-faire outcome where only susceptible people have an incentive to do so. Yet a social planner need not be restricted in this manner and may choose to curtail the movement of all people, regardless of disease status. It's clear that our central insights carry over to this case, but with quantitative differences.

Second, we only consider travel restriction policies which are fixed throughout the epidemic. In practice, governments do vary these over time as the epidemic and economic conditions change and it would be interesting to explicitly analyse how to best go about this. During COVID-19, it was clear that travel restrictions were not reciprocal or even coordinated between countries, raising interesting questions of strategic behaviour in the imposition of these over time.

Third, our analysis is based on the assumption that the underlying disease dynamics are of the SIR type. But it is easily extended to richer environments, such as the more

general SEIRS model with pre-symptomatic infection and waning immunity. Again, many of our conclusions will carry over to such a setting.

Fourth, our mitigation policy of choice can best be interpreted as measures such as lockdowns or mandated social distancing, yet the insights would be valid also with other control measures such as mass vaccination. An analysis of international vaccine rollout, using our methodology, would speak to recent policy debates around vaccine nationalism and global vaccine inequalities.

Fourth, as noted by Ndeffo Mbah and Gilligan (2011), implementing non-uniform policies may be considered inequitable and may cause local resistance to mitigation efforts even if they achieve desirable aggregate outcomes. This was the case in Manchester, UK in late 2020, where political leaders requested compensation from the central government in return for cooperation in implementing robust "Tier 4" lockdown measures.[15] This raises the prospect of using inter-regional or even cross-country transfers to provide incentives for disease mitigation.

Last, our analysis uses a relatively simple model of two symmetric regions. Expanding the model to a network of many asymmetrically coupled regions may be more computationally complex, but would allow for richer spatio-temporal dynamics of disease spread. This seems an avenue worth pursuing in future work.

References

Acemoglu, D., Chernozhukov, V., Werning, I., and Whinston, M. D. (2021). Optimal targeted lockdowns in a multigroup sir model. *American Economic Review: Insights*, *3*(4), 487–502.

Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2020). Testing, voluntary social distancing and the spread of an infection. *NBER Working Paper Series No. 27483*.

Birge, J. R., Candogan, O., and Feng, Y. (2022). Controlling epidemic spread: Reducing economic losses with targeted closures. *Management Science*, *68*(5), 3175–3195.

Bognanni, M., Hanley, D., Kolliner, D., and Mitman, K. (2020). Economics and epidemics: Evidence from an estimated spatial econ-sir model. *Finance and Economics Discussion Series 2020-091*.

Byrne, A. W., McEvoy, D., Collins, A. B., Hunt, K., Casey, M., Barber, A., ... More, S. J. (2020). Inferred duration of infectious period of sars-cov-2: rapid scoping review and analysis of available evidence for asymptomatic and symptomatic covid-19 cases. *BMJ Open*, *10*(8).

Caputo, M. (2005). *Foundations of dynamic economic analysis: Optimal control theory and applications*. Cambridge University Press.

Chen, F. (2012). A mathematical analysis of public avoidance behavior during epidemics using game theory. *Journal of Theoretical Biology*, *302*, 18–28.

Chen, F., Jiang, M., Rabidoux, S., and Robinson, S. (2011). Public avoidance and epidemics: Insights from an economic model. *Journal of Theoretical Biology*, *278*, 107–119.

Dockner, E., Jorgensen, S., Long, N., and Sorger, G. (2000). *Differential games in economics and management science*. Cambridge University Press.

Fajgelbaum, P., Khandelwal, A., Kim, W., Mantovani, C., and Schaal, E. (2021). Optimal lockdown in a commuting network. *American Economic Review: Insights*, *3*(4), 503–522.

Farboodi, M., Jarosch, G., and Shimer, R. (2021). Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory*, *196*, 105293.

Fenichel, E. P. (2013). Economic considerations for social distancing and behavioral based policies during an epidemic. *Journal of Health Economics*, *32*(2), 440–451.

Keeling, M., and Rohani, P. (2008). *Modeling infectious diseases in humans and animals*. Princeton University Press.

Kucharski, A. J., Russell, T. W., Diamond, C., Liu, Y., Edmunds, J., Funk, S., and Eggo, R. M. (2020). Early dynamics of transmission and control of covid-19: a mathematical modelling study. *The Lancet Infectious Diseases*, *20*(5), 553–558.

Makris, M. (2021). Covid and social distancing with a heterogenous population. *Economic Theory*.

Makris, M., and Toxvaerd, F. (2020). Great expectations: Social distancing in anticipation of pharmaceutical innovations. *Covid Economics*(56), 1–19.

Ndeffo Mbah, M. L., and Gilligan, C. A. (2011). Resource allocation for epidemic control in metapopulations. *PLoS ONE*, *6*(9).

Reluga, T. C. (2010). Game theory of social distancing in response to an epidemic. *PLoS Computational Biology*, *6*(5).

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

Rowthorn, R. E., and Toxvaerd, F. (2020). The optimal control of infectious diseases via prevention and treatment. *Cambridge Working Papers in Economics CWPE2027*.

Seierstad, A., and Sydsaeter, K. (1987). *Optimal control theory with economic applications* (Vol. 24). North Holland.

Sethi, S. P. (1978). Optimal quarantine programmes for controlling an epidemic spread. *The Journal of the Operational Research Society*, *29*(3), 265–268.

Toxvaerd, F. (2019). Rational disinhibition and externalities in prevention. *International Economic Review*, *60*(4), 1737–1755.

Toxvaerd, F. (2020). Equilibrium social distancing. *Cambridge Working Papers in Economics CWPE2021*.

Wu, D., Wu, T., Liu, Q., and Yang, Z. (2020). The sars-cov-2 outbreak: What we know. *International Journal of Infectious Diseases*, *94*, 44–48.

Yarmol-Matusiak, E. A., Cipriano, L. E., and Stranges, S. (2021). A comparison of covid-19 epidemiological indicators in sweden, norway, denmark, and finland. *Scandinavian Journal of Public Health*, *49*(1), 69–78.

A Transversality conditions

The current-value costate variable $\lambda_{\tau_j}(t)$ has the economic interpretation of capturing the shadow value of the disease state τ_j at time t (Caputo, 2005). As such, the transversality conditions of $\lambda_{\tau_j}(t)$ in the case of a *finite* time-horizon ending at time T are

$$\lambda_{S_j}(T) = V_S \tag{A.1a}$$

$$\lambda_{I_j}(T) = V_I \tag{A.1b}$$

$$\lambda_{R_j}(T) = V_R \tag{A.1c}$$

where V_τ is the expected present value of utility earned (discounted to time T) by an individual in disease compartment τ_j at time T.

In the case of the *infinite* time-horizons considered in this paper, the corresponding transversality conditions are

$$\lim_{T\to\infty} \lambda_{S_j}(T) = V_S \tag{A.2a}$$

$$\lim_{T\to\infty} \lambda_{I_j}(T) = V_I \tag{A.2b}$$

$$\lim_{T\to\infty} \lambda_{R_j}(T) = V_R \tag{A.2c}$$

For V_R, an individual who is in the recovered compartment at time T will remain in the recovered compartment for the rest of time. Thus, there is no uncertainty about their future disease state or utility flows. V_R can therefore be calculated as

$$V_R = \int_T^\infty e^{-\rho(t-T)}\pi_{\mathcal{R}}dt = \pi_{\mathcal{R}}\int_T^\infty e^{-\rho(t-T)}dt = \frac{\pi_{\mathcal{R}}}{\rho} \tag{A.3}$$

For individuals in the infected compartment at time T, their recovery and transition to the recovered compartment occurs according to a Poisson process with rate $\gamma > 0$. Let δ denote the amount of time after T at which the individual transitions to the recovered compartment. The probability distribution function of δ is therefore $f_\delta(t) = \gamma e^{-\gamma t}$. In the period $t \in [T, T+\delta)$, the individual earns utility flows of $\pi_{\mathcal{I}}$. Once transitioning at

time $t = T + \delta$, the individual earns a net present value of V_R. Thus, V_I is calculated as

$$\begin{aligned}
V_I &= \int_0^\infty f_\delta(t) \left[\frac{\pi_{\mathcal{I}}}{\rho} + e^{-\rho t} \left(\frac{\pi_{\mathcal{R}}}{\rho} - \frac{\pi_{\mathcal{I}}}{\rho} \right) \right] dt \\
&= \int_0^\infty \frac{\gamma}{\rho} e^{-\gamma t} \left[\pi_I + e^{-\rho t} \left(\pi_{\mathcal{R}} - \pi_{\mathcal{I}} \right) \right] dt \\
&= \frac{\gamma}{\rho} \int_0^\infty e^{-\gamma t} \pi_{\mathcal{I}} dt + \frac{\gamma}{\rho} \int_0^\infty e^{-(\gamma+\rho)t} \left(\pi_{\mathcal{R}} - \pi_{\mathcal{I}} \right) dt \\
&= \frac{\pi_{\mathcal{I}}}{\rho} + \frac{\gamma}{\rho} \frac{\pi_{\mathcal{R}}}{\gamma + \rho} - \frac{\gamma}{\rho} \frac{\pi_{\mathcal{I}}}{\gamma + \rho} \\
V_I &= \frac{1}{\gamma + \rho} \left(\pi_{\mathcal{I}} + \frac{\gamma}{\rho} \pi_{\mathcal{R}} \right)
\end{aligned} \tag{A.4}$$

Lastly, individuals who are susceptible at time T have a chance of incurring the infection at some time after T, which depends on the infection levels and their dynamics beyond T. However, as $\pi_{\mathcal{S}} \geq \pi_{\mathcal{R}} \geq \pi_{\mathcal{I}}$, V_S is bounded between V_I and $\frac{\pi_{\mathcal{S}}}{\rho}$. The former assumes that the individual immediately becomes infected after T, whereas the latter assumes the individual never incurs the infection.

As an interior solution to the mitigation problem is assumed ($m_j(t) \in (0, 1),\ \forall t$), complete mitigation never takes place. As a result, there will always be at least some interaction between susceptibles and infecteds and new infections will occur. The epidemic will therefore always continue to evolve under endogenous mitigation, and will never be halted. Thus, in the infinite horizon, herd immunity will eventually be achieved and the infection prevalence will asymptotically disappear. In other words, if T is sufficiently large, infections will be effectively eradicated and $V_S \approx \frac{\pi_{\mathcal{S}}}{\rho}$. In the limit, we have

$$\lim_{T \to \infty} V_S = \frac{\pi_{\mathcal{S}}}{\rho} \tag{A.5}$$

A.1 Delegated policy setting's costate variables

The current-value costate variables $\lambda^j_{\tau_k}(t)$ in the regional social planner's Hamiltonian (4.14) are unique in that they depend not only on the disease state τ_k, but also on the region j of the regional social planner. Recall that $\lambda^j_{\tau_k}(t)$ captures the shadow value of τ_k to regional social planner j. As the payoffs earned in the opposing region are not valued by the regional social planner, $\lambda^j_{\tau_{-j}}(t)$'s shadow value is solely based on disease state τ_{-j}'s implied impact on the disease dynamics and payoffs experienced in region j. This is similarly reflected in the transversality conditions. For the finite time-horizon

approximation, we have

$$\lambda^j_{S_j}(T) = V_S \tag{A.6a}$$
$$\lambda^j_{I_j}(T) = V_I \tag{A.6b}$$
$$\lambda^j_{R_j}(T) = V_R \tag{A.6c}$$
$$\lambda^j_{S_{-j}}(T) = \lambda^j_{I_{-j}}(T) = \lambda^j_{R_{-j}}(T) = 0 \tag{A.6d}$$

In the infinite time-horizon, the corresponding transversality conditions are

$$\lim_{T\to\infty} \lambda^j_{S_j}(T) = V_S \tag{A.7a}$$
$$\lim_{T\to\infty} \lambda^j_{I_j}(T) = V_I \tag{A.7b}$$
$$\lim_{T\to\infty} \lambda^j_{R_j}(T) = V_R \tag{A.7c}$$
$$\lim_{T\to\infty} \lambda^j_{S_{-j}}(T) = \lim_{T\to\infty} \lambda^j_{I_{-j}}(T) = \lim_{T\to\infty} \lambda^j_{R_{-j}}(T) = 0 \tag{A.7d}$$

B Impact of increased mitigation in region 2

To further assess the impact that mitigation levels in region 2 have on mitigation levels in region 1, we consider an extension to the quasi-decentralised setting in which region 2's mitigation is set at a constant level $y \in [0, 1]$. Meanwhile, mitigation in region 1 is still determined by laissez-faire. Let $\hat{m}^y_1(t)$ denote the corresponding equilibrium mitigation level of susceptible individuals in region 1. The disease dynamics in this scenario are then described by the metapopulation SIR model with mitigation (2.4) where $\bar{m}_1(t) = \hat{m}^y_1(t)$ and $\bar{m}_2(t) = y$.

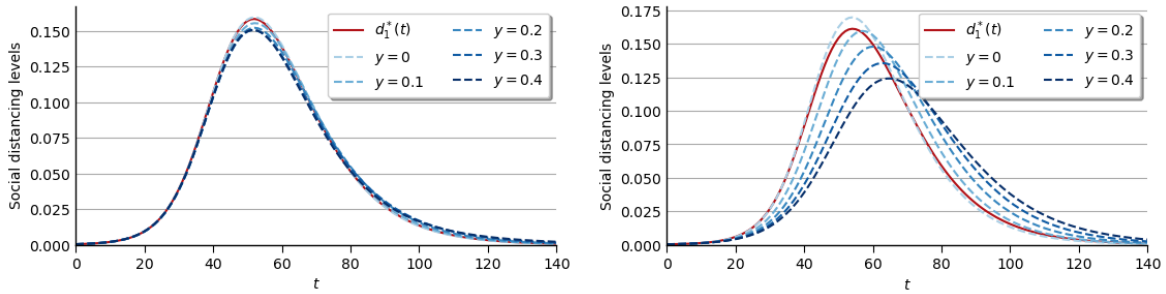

(a) Low ε. (b) High ε.

Figure 28: Equilibrium mitigation levels in region 1 for various mitigation levels in region 2 ($\bar{m}_2(t) = y$) and in the laissez-faire setting ($m^*_1(t)$).

Figures 28 and 29 show that when mitigation levels in region 2 increase, the impact on mitigation levels in region 1 changes over the course of the pandemic. Namely, it induces individuals to do less mitigation in the early stages of the epidemic, but more

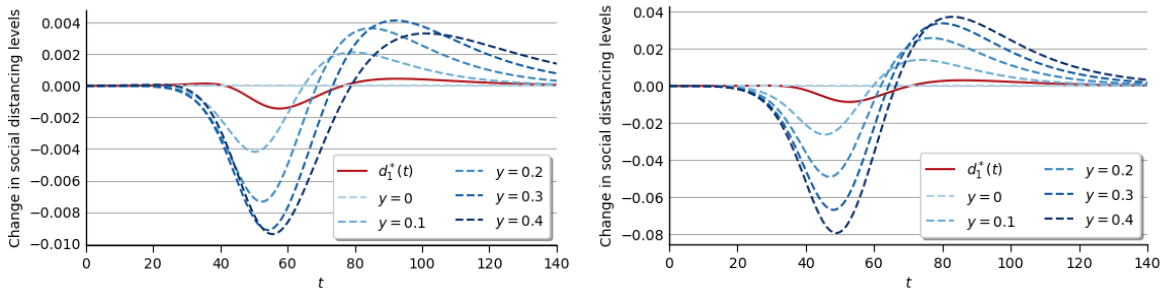

(a) Low ε. (b) High ε.

Figure 29: Absolute change in mitigation levels in region 1 for various mitigation levels in region 2 ($m_2(t) = y$) and under decentralised decision making ($m_j^*(t)$), relative to $y = 0$.

mitigation towards the end. Furthermore, the impact is greater when regions are more strongly connected.

C Numerical solution methods

To obtain many of our insights, we rely on numerical solutions of the metapopulation SIR model, both without mitigation (2.1) and with mitigation (2.4). This section details the general steps we undertake to generate the simulations.

We solve the model numerically with Python using the *odeint*() and *solve_bvp*() functions from the 'integrate' module of SciPy. These functions require as inputs a system of ordinary differential equations and boundary conditions, and returns a solution to the ODEs. For the boundary conditions, the *odeint*() function only allows for initial value conditions, whereas the *solve_bvp*() function also allows for terminal value conditions. Default options for each of these functions were used, with the exception of a stricter tolerance level for the *solve_bvp*() function of $tol = 1e - 5$.

To initialise the process, we create a time grid containing points for each day of $t \in [0, T]$, resulting in a vector of size $T + 1$. This time grid is what we lay the solutions to our ODEs on top of.

Uncontrolled epidemiological benchmark

First, we solve for the mechanistic epidemiological model without mitigation (2.1). For this, we use the *odeint*() function with the following steps:

1. We create an initial guess for each disease state vector $[\tau_j(0), \tau_j(1), \ldots, \tau_j(T)]$, for which we set all elements to the initial value $\tau_j(0)$. The initial values we use are:

$$S_1(0) = 0.999, \quad I_1(0) = 0.001, \quad R_1(0) = 0, \tag{C.1a}$$

$$S_2(0) = 1, \quad I_2(0) = R_2(0) = 0 \tag{C.1b}$$

2. Using the initial values of the disease compartments (C.1), along with the differential equations from (2.1), we run the *odeint*() function to obtain the value of each regional disease compartment's proportion along the time grid.

Endogenous mitigation

For the laissez-faire, targeted global social planner, and uniform policy settings, we use the *solve_bvp*() function to solve for 12 time-varying variables: 6 disease compartment variables ($\tau_j(t)$) and 6 current-value costate variables ($\lambda_{\tau_j}(t)$). In the delegated policy setting, we need to solve for 18 variables, as there are now 6 costate variables for each region ($\lambda^j_{\tau_k}(t)$). Note that the mitigation levels do not need to be solved for by the function, as they are calculated implicitly according to (3.6), (4.6), (4.11), and (4.16). We complete this process with the following steps

1. First, we create an initial guess for each of the variables.
   (a) For the disease compartments $\tau_j(t)$, we use the previous solutions from the uncontrolled epidemiological benchmark as the initial guess.
   (b) For the costate variables $\lambda_{\tau_j}(t)$, we use their terminal values, as given in (A.1). In the case of the delegated policy setting's costate variables $\lambda^j_{\tau_k}(t)$, we use their corresponding terminal values, as given in (A.6).
2. We then define the differential equations for the variables. The disease state variables' differential equations are given by the metapopulation model with mitigation (2.4), in which the mitigation level is calculated as the solution corresponding to the setting that we are simulating (3.6), (4.6), (4.11), (4.16). Similarly, the costate variables' differential equations are defined as per the simulation's setting, given in (3.4), (4.5), (4.10), (4.15).
3. We specify the variables' boundary conditions. In particular, the initial value conditions of the disease states $\tau_j(0)$ and the terminal value conditions of the costate variables ($\lambda_{\tau_j}(t), \lambda^j_{\tau_k}(t)$), as detailed in Step 1.
4. Using these differential equations and boundary conditions, we run the *solve_bvp*() function to obtain the solutions for each of the time-varying variables under the setting we are simulating. We then use these solutions to calculate the vector of mitigation levels along the time grid.

After obtaining the results of the simulation, we perform a brief sanity check. Namely, we ensure that the calculated mitigation levels indeed correspond to an interior solution ($m_j(t) \in (0,1)$) and that each of the disease compartment variables have feasible values ($\tau_j(t) \in [0,1]$). Given the approximate nature of numerical solution methods, we allow for small deviations of up to $1e-7$ beyond these bounds.

As discussed in Appendix A, we need to approximate the infinite horizon with a long finite horizon of T. For our simulations, we use $T = 200$. To ensure this is a reasonable approximation, we check the infection prevalences at the terminal date to ensure they are sufficiently close to 0 (i.e., $I_j(T) < 1e-4$). Furthermore, we found that using larger values of T did not noticeably change our results.

For the disease dynamics parameters, we use $R_0 = 2.25,\ \gamma = 1/7.5,$ and (a) low $\varepsilon = 0.025$ (b) high $\varepsilon = 0.25$, as discussed in Section 2.3. Additionally, we use a daily discount rate of $\rho = 0.05/365$, which corresponds to a 5% annual interest rate. Finally, for the payoff parameters, we use $\pi_{\mathcal{S}} = \pi_{\mathcal{R}} = 0$, $\pi_{\mathcal{I}} = -1$, and $c = 1$.

D WTP dependence on ρ

D.1 Without mitigation

To assess whether the WTP to decrease ε in the case of no mitigation is purely driven by the change in epidemic timings, we calculate and compare the WTP using different discount rates ρ. The rest of the parameter values are the same as before, namely: $I_1(0) = 0.001,\ I_2(0) = 0,\ R_1(0) = R_2(0) = 0,\ R_0 = 2.25,\ \gamma = 1/7.5,\ \pi_{\mathcal{S}} = \pi_{\mathcal{R}} = 0,$ $\pi_{\mathcal{I}} = -1$ and $c = 1$.

As $\rho \to 0$, individuals become more patient and place less value on the timing of utility flows. They therefore become increasingly indifferent towards an epidemic's timing, which is what is altered when ε changes. As can be seen in Figure 30, WTP$(\hat{\varepsilon})$ becomes negligible as $\rho \to 0$.[16] This corroborates the idea that the WTP to decrease ε in the case of no mitigation is purely a timing issue.

D.2 With endogenous mitigation

Figures 31–33 depict the WTP curves under the laissez-faire, targeted global social planner, and delegated policy settings, respectively. In all cases, the WTP curves' non-monotonicity caused by the inclusion of endogenous mitigation becomes more exaggerated and apparent as $\rho \to 0$. This is due to the fact that the lengthened epidemic durations have a more noticeable impact through the infection differential effect when ρ is low. Thus, the differences between the WTP in the mechanistic case and that under endogenous mitigation are more pronounced for lower values of ρ.

[16] This is true except for the unique case of $\hat{\varepsilon} = 0.001$, in which decreasing ε to 0 entails entirely disconnecting the two regions, such that the initial infection in region 1 cannot spread to region 2. Hence, the calculated WTP in this case is the WTP to avoid the epidemic altogether in region 2.

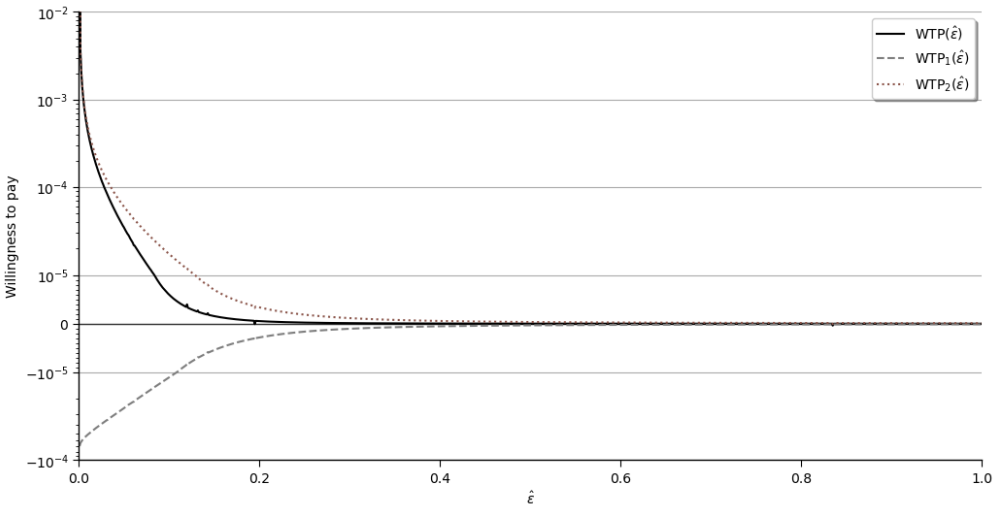

(a) $\rho = 0.05/365$.

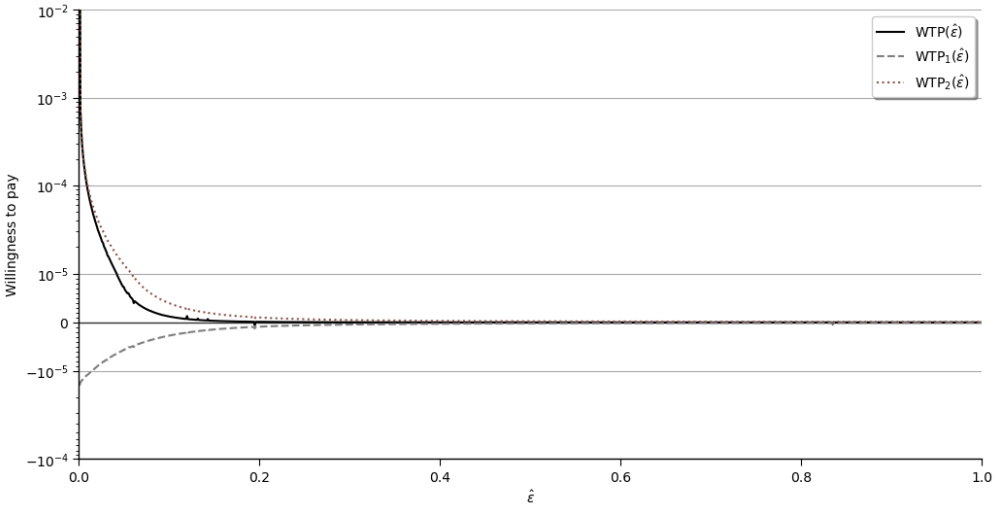

(b) $\rho = 0.01/365$.

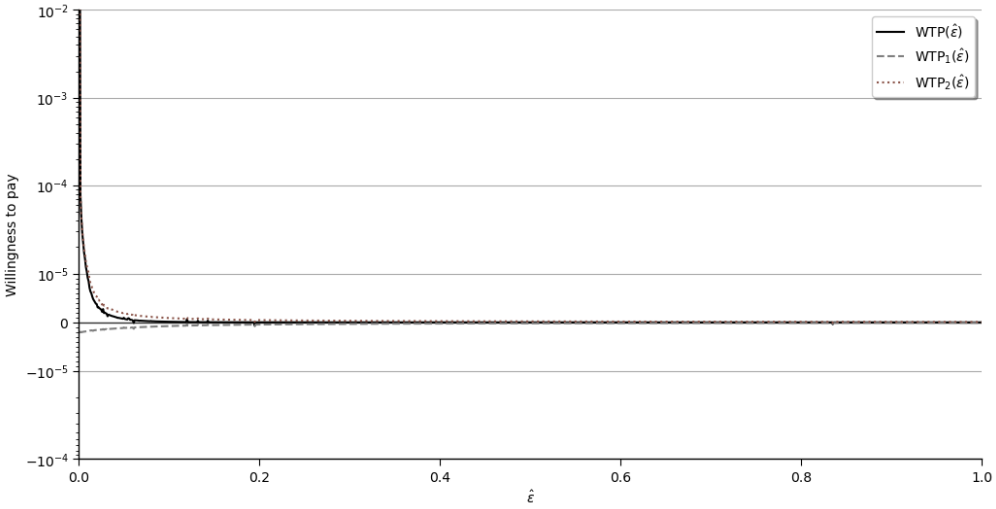

(c) $\rho = 0.001/365$.

Figure 30: Regional and total WTP to decrease ε by 0.001 in the mechanistic model with no mitigation, using different discount rates ρ.

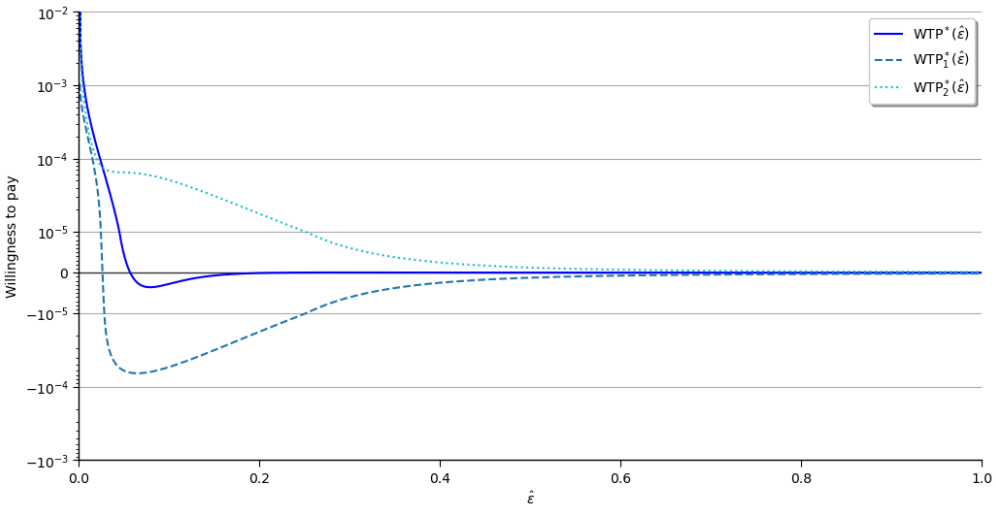

(a) $\rho = 0.05/365$.

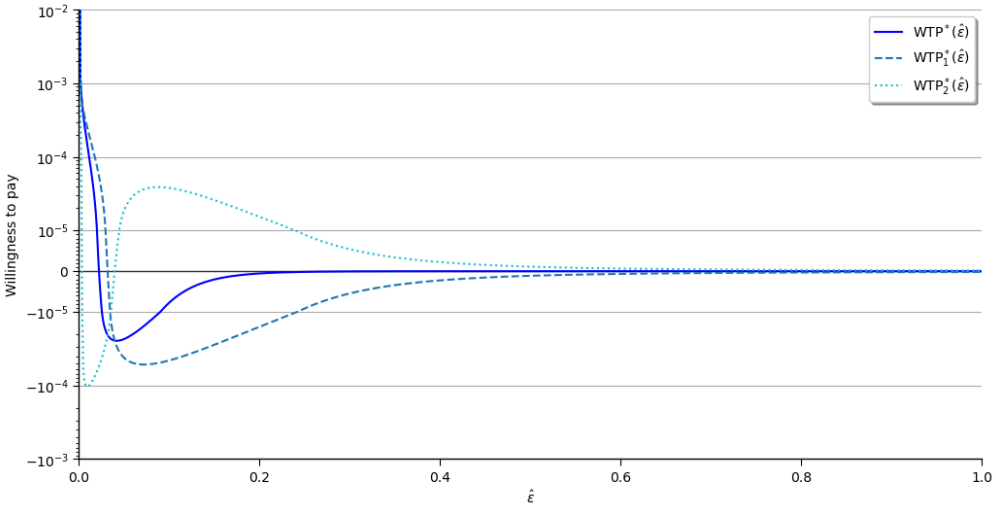

(b) $\rho = 0.01/365$.

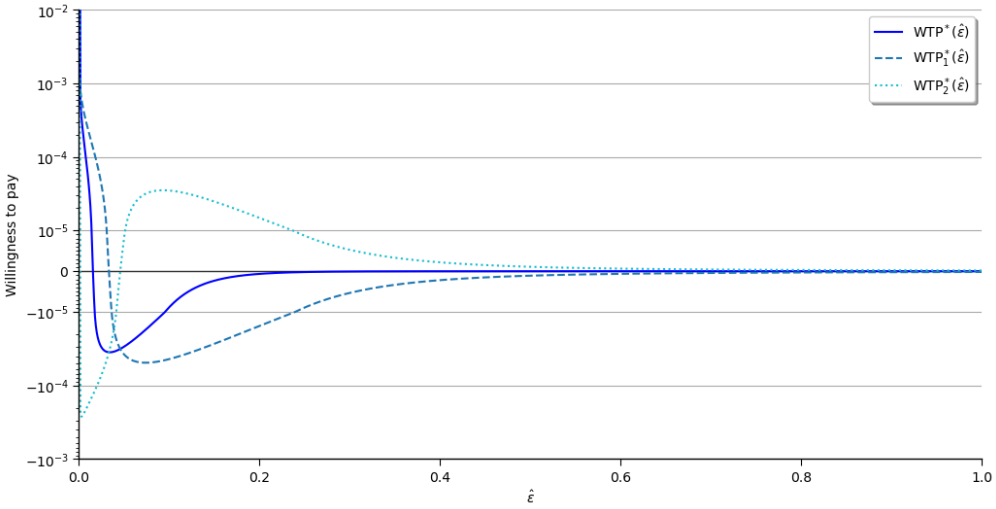

(c) $\rho = 0.001/365$.

Figure 31: Regional and total WTP to decrease ε by 0.001 under laissez-faire mitigation, using different discount rates ρ.

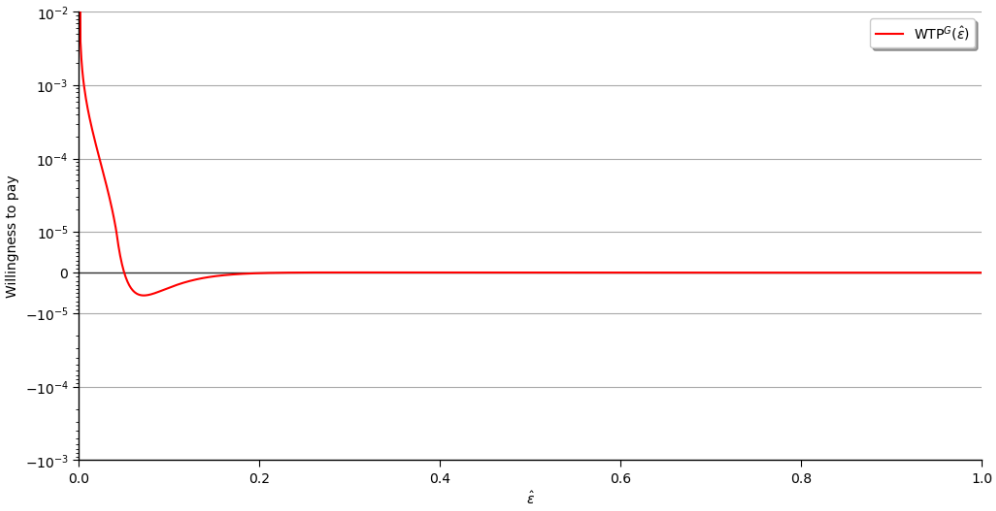

(a) $\rho = 0.05/365$.

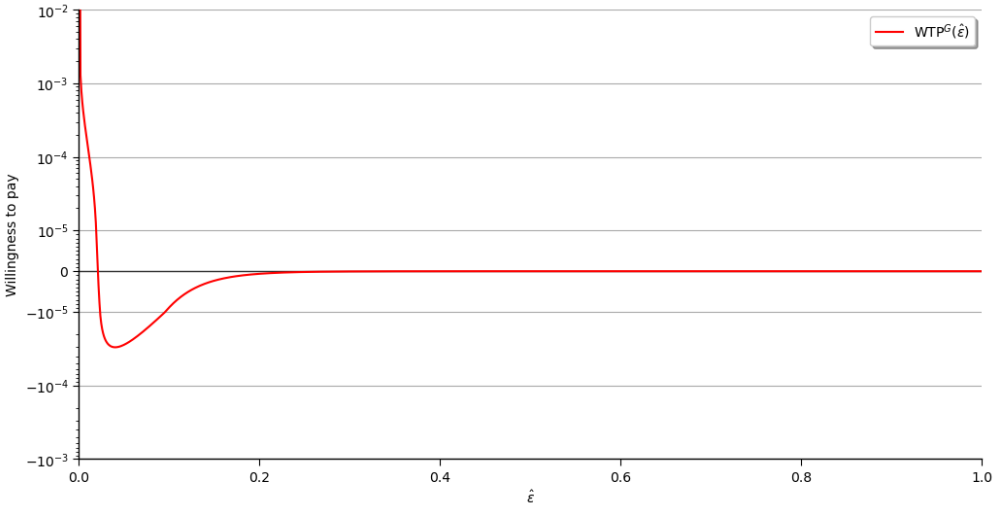

(b) $\rho = 0.01/365$.

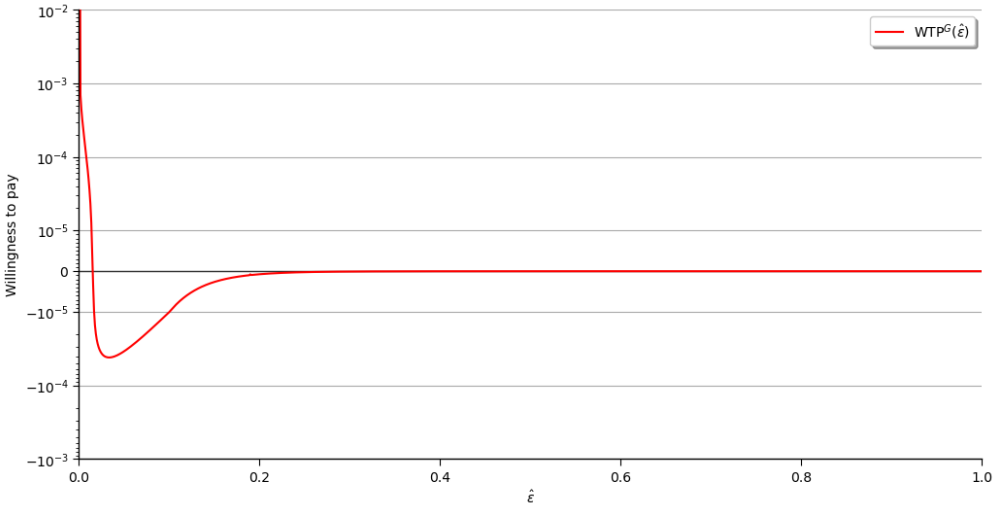

(c) $\rho = 0.001/365$.

Figure 32: Total WTP to decrease ε by 0.001 in the targeted global social planner setting, using different discount rates ρ.

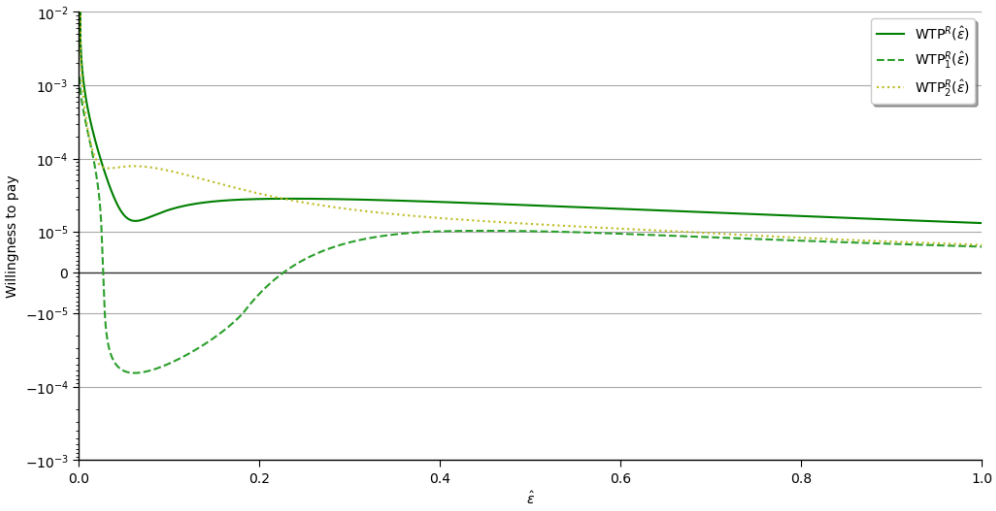

(a) $\rho = 0.05/365$.

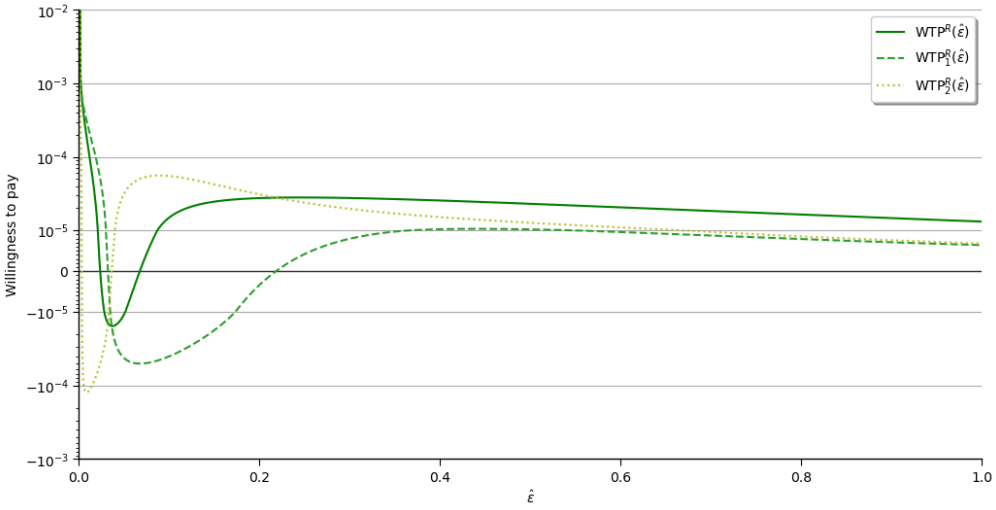

(b) $\rho = 0.01/365$.

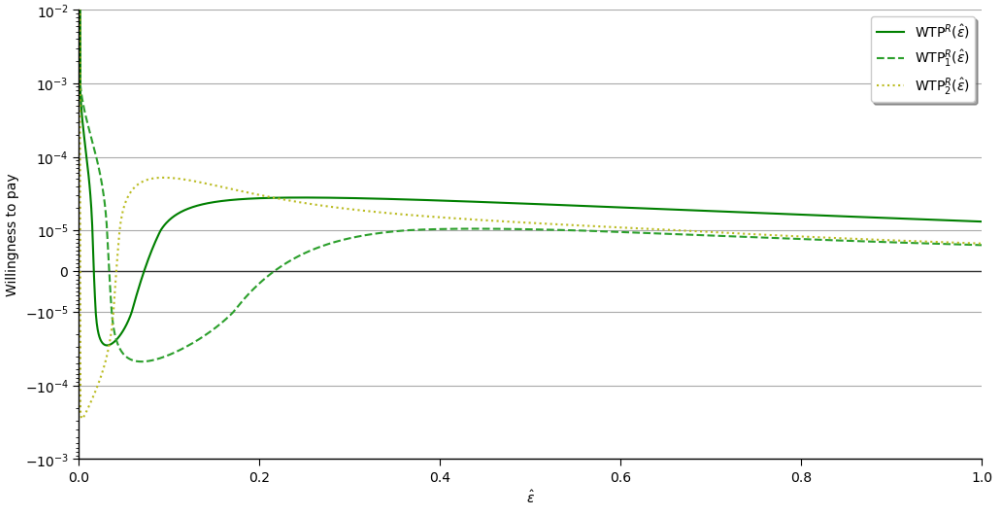

(c) $\rho = 0.001/365$.

Figure 33: Regional and total WTP to decrease ε by 0.001 in the delegated policy setting, using different discount rates ρ.