Pandemic Policy in a Global Game

*For the most recent version, click **here***

Roweno J.R.K. Heijmans and Ana Moura*

February 11, 2023

Abstract

We study pandemic policy in a global game. Independent regions face the outbreak of a disease and can exert effort to control it. If the disease is infectious, it will be controlled only if sufficiently many regions exert effort. Regions thus face a coordination game which, in the well-studied case of perfect information, has multiple Nash equilibria. We show that even a vanishing amount of uncertainty about fundamentals of the game forces selection of a unique equilibrium. In well-identified cases, a pandemic occurs even though it is inefficient and could be avoided. Harmful diseases are less likely to become an pandemic while diseases which require greater cooperation have a larger chance to go uncontrolled. A higher cost of effort increases the probability of a pandemic and regions whose actions exhibit stronger spillovers have a greater influence on the limiting equilibrium. Given the possibility of an inefficient but rational pandemic, we also propose a novel mechanism to facilitate coordination on disease control. In particular, we introduce the concept of P-delegation, a kind of conditional delegation which binds delegating regions to exert effort toward disease control if and only if at least a weighted proportion P of all regions has delegated. We show that, for judiciously chosen P, P-delegation can help avoid pandemics. Our results thus suggest a way forward for the international cooperation on disease control.

Keywords: global games, pandemics, disease eradication, privately provided public goods

JEL Codes: I18, H41, C72

1 Introduction

Pandemics are costly. Though intuition might attribute the rise of a pandemic to mere misfortune, that explanation leaves open the question how individual decisions affect the likelihood of a pandemic's onset. This paper uses game theory to study the strategic

*Heijmans: Department of Economics, Swedish University of Agricultural Sciences, Uppsala, Sweden (e-mail: roweno.heijmans@slu.se). Moura: Modeling and Meta-Analysis team, OPEN Health, Rotterdam, the Netherlands (e-mail: ana.cdmcm@gmail.com). We gratefully acknowledge Reyer Gerlagh, Jan Boone, Jeffrey Campbell, and Sjak Smulders for invaluable comments.

considerations leading up to a pandemic. Our game of imperfect information is distinct from existing studies in the economics of epidemics that either do not use game theory (Kremer, 1996; Geoffard and Philipson, 1997; Gersovitz and Hammer, 2003), or model the problem as a game of perfect information (Barrett, 2003). We formulate several novel insights.

Game theory provides a tool to analyze strategic interactions between rational decision-makers. It captures the fundamentally strategic nature of disease control problems, where any one region's optimal course of action may well depend on what other regions are doing. However, extant game theoretic analyses of disease control, see Barrett (2003) in particular, tend to have multiple equilibria. Equilibrium multiplicity dwindles the predictive ability of these models. The outcome of a game with multiple equilibria is not, a priori, determined – there is no apparent connection between a disease's fundamental properties and its eventual fate. In such games, a pandemic is mere misfortune indeed.

The indeterminacy of outcomes in existing models comes about through a combination of perfect information and the possibility of contagion. An epidemic spreads from one region to another. Any region's private efforts at controlling the spread of a disease thus boost the likelihood that other regions' private efforts are successful. This type of mutual reinforcement, when combined with perfect information, often leads to a manifold of equilibria and coordination failure (*c.f.* Van Huyck et al. (1990)).

While our game maintains the strategic nature of disease control, we study disease control under payoff uncertainty, which makes for a global game (Carlsson and van Damme, 1993; Frankel et al., 2003). Global games are a class of imperfect information games where players are uncertain about some underlying fundamental of the game but receive private noisy signals of it. In our model, regions are uncertain about the precise number of regions that need to exert effort in order to successfully control the disease.

There are at least two reasons to consider this type of uncertainty. First, pandemics have a highly multi-dimensional (negative) impact on society, making a comprehensive idea of their costs hard to grasp. The costs extend far beyond the mere expenditures on health care or illness-related loss of productivity (Shastry and Weil, 2003). They can lead to civil conflict (Cervellati et al., 2017, 2018) and cause impaired development in youth (Bleakley, 2003; Coffey et al., 2017). An pandemic may call the government's legitimacy into question (Flückiger and Ludwig, 2019) and can even affect a country's institutions (Acemoglu et al., 2001, 2003). Second, the epidemiological literature suggests that new diseases – by their nature subject to many uncertainties – will appear increasingly often in the near future (see Rappuoli, 2004, for a comprehensive review). A complete understanding of disease control under incomplete information is therefore necessary to inform future health policy.

Even a vanishing amount of uncertainty leads to selection of a unique equilibrium in our disease control game. While equilibrium uniqueness is well-established for global games generally, see especially Frankel et al. (2003), it is a new insight in the literature on pandemics. The advantage of a unique equilibrium over multiple equilibria is that it allows for sharper predictions. In our global game, a more harmful disease is less likely to become an pandemic. In terms of welfare, an pandemic can occur even though this is inefficient and could have been avoided. These results embed well-known predictions from the epidemiological studies within a strategic framework (*c.f.* Capasso and Serio (1978); Epstein et al. (2008); Fenichel et al. (2011)) and are supported empirically (Ahituv et al., 1996; Philipson, 1996). Also intuitive, a disease that is harder to control is more likely to become a pandemic, higher costs of effort increase the chances of a pandemic, and regions whose effort has stronger spillovers have a bigger impact on the limiting equilibrium of the game.

The result that regions may rationally coordinate on a pandemic calls for a coordination mechanism. Finding a suitable coordination device is not straightforward because sovereign regions (e.g. countries) cannot be coerced into a course of action they do not themselves wish to take. Because of this, we take inspiration from the literature on strategic delegation and allow regions to delegate a third-party authority to take binding decisions on the their behalf. Strategic delegation by countries has been studied in a variety of contexts including trade policy in customs unions (Gatsios and Karp, 1991), international climate policy (Habla and Winkler, 2018), and the fight against terrorism (Siqueira and Sandler, 2007). However, simple strategic delegation may not work in this context as the decision to delegate one's choice of effort can itself be thought of as a coordination problem. To resolve this, we introduce the novel concept of P-delegation; like standard delegation, P-delegation lets the regions delegate their decision-making powers to a third agents. However, the decision made by the agent is binding if and only if sufficiently many other regions also delegated their decisions. We show that P-delegation, given a judiciously chosen minimum participation threshold P solves the global coordination problem and can thus help avoid pandemics.

The remainder of this paper is organized as follows. Section 2 presents the building blocks of our model and the main results. Section 3 incorporates cooperation into the model. Section 4 concludes. The proofs of our main results are in the text and the proofs of additional results are in the Appendix.

2 Model

There is a set $N = \{1, 2, ..., n\}$ of sovereign regions, each plagued by the same infectious disease. We follow Barrett (2003) and assume that each region i can either exert effort to control the spread of the disease ($x_i = 1$), or not ($x_i = 0$), and regions choose their actions simultaneously. The term "effort" here is understood to be an umbrella term describing any attempt toward containment or eradication; what precisely it entails may depend on the disease in question. We let $x = (x_i)$ denote the effort vector of all regions and $X = \sum_{i \in N} \omega_i x_i / n$, where $\sum_i \omega_i = 1$ and $\omega_i > 0$ for all i. The parameter ω_i is a measure for the importance of region i's eradication efforts in a way we make precise below. The cost of effort to region i is denoted $c_i > 0$. If a region succeeds in controlling the disease, it realizes a benefit b_i and $b_i > c_i$. We assume that there is a critical mass $T \in [0, 1]$ of regions such that effort is successful if and only if $X \geq T$. This captures the idea that contagious or infectious diseases can easily spread from one region to another so that some amount of coordinated effort is required to stand a reasonable chance toward successful control. Note that a disease is, for all practical purposes, fully characterized by the tuple (b, c, T), where $b = (b_i)_{i \in N}$ and $c = (c_i)_{i \in N}$ are vectors of benefit and cost parameters, respectively, for all regions. We interpret the parameter T as a measure of how difficult it is to control a disease.

The payoff to region i is given by

$$u_i(x; T) = \begin{cases} b_i - c_i & \text{if} \quad X \geq T, \\ -c_i & \text{if} \quad X < T, \end{cases} \tag{1}$$

where we normalize the payoff to no effort ($x_i = 0$) to 0.[1]

The payoff function (1) makes clear that regions have an incentive to coordinate their efforts. To keep the example interesting, assume that $\omega_i < T$ for all i. Then each region wants to exert effort, provided the critical threshold for success T is passed. If however all regions i believe that $X < T$ (even counting their own effort), none want to exert effort. The game therefore has two pure strategy Nash equilibria, one in which regions coordinate on controlling the disease, another in which regions coordinate on *not* controlling the disease.

Observation 1. *In the game of perfect information about (b, c, T), if $b_i > c_i$ and $\omega_i < T$ for all i, there are two pure strategy Nash equilibria, one in which $x_i = 1$ for all i, another in which $x_i = 0$ for all i.*

Observation 1 establishes equilibrium multiplicity in the perfect information game and is equivalent to Barrett's (2003) Proposition 3. Observe that a situation in which only a subset of regions exerts effort cannot be an equilibrium, even when regions are asymmetric. Once the critical threshold T is passed, regions that do not exert effort have an incentive to change course and control the disease. Similarly, if the critical threshold T is not passed, those regions that do exert effort have an incentive to not do so and save on the cost (unsuccessful) effort.

It would seem a reasonable proposition that a multiplicity of equilibria per se is no reason for concern. As Barrett (2003) observes, in the presence of multiple equilibria, "presumably, [... regions] will try to coordinate on the preferred equilibrium." Unfortunately, as we will shortly show, the co-existence of two Pareto-ranked equilibria (such that regions might 'choose' the one they prefer) is an artifact of the extreme assumption of complete and perfect information about the game. Even minimal uncertainty about the game played leads to selection of a unique equilibrium that generally is not the Pareto-dominant outcome of the game. A commitment mechanism will be needed to allow regions to always coordinate on the efficient outcome of the game; we discuss one possible such mechanism in Section 4.

3 Equilibrium selection and rational pandemics

The assumption that regions have perfect and common knowledge of the entire game is strong. For the case of an infectious disease, it may be too strong. We therefore discuss the case of incomplete information in this section.

Suppose that regions do not know exactly what the critical threshold T for successful control is; that is, regions are uncertain how difficult it actually is to control the disease. Let it be common knowledge that T is drawn from the uniform distribution on $[0, 1]$. Each region i receives a private noisy signal t_i^η of T, given by

$$t_i^\eta = T + \eta \cdot \varepsilon_i,$$

[1] The model can accommodate free-riding. Suppose a region i that does not exert effort can nevertheless benefit from successful control ($X \geq T$); let these no-effort benefits be denoted $\underline{b}_i$. The benefit of successful control conditional on effort is $\overline{b}_i$. If we assume that $\overline{b}_i > \underline{b}_i$, for example because more people will get sick prior to eventual control in regions that do not exert effort, we can define $b_i := \overline{b}_i - \underline{b}_i$ and use the payoff function u_i specified in (1).

where $\eta \in (0,1]$ is a scaling factor for the random error term ε_i. The distribution function of ε_i is common knowledge and given by F_i with support $[-1/2, 1/2]$ and continuous density f_i. Observe that the distribution of t_i^ε conditional on T is given by $F_i\left(\frac{t_i^\varepsilon - T}{\eta}\right)$. Let $\mathcal{T} = [-\eta/2, 1+\eta/2]$ denote the set of feasible signals a region i can receive.

A strategy $s_i : \mathcal{T} \to [0,1]$ for region i is a function that assigns to any feasible signal $t_i^\eta \in \mathcal{T}$ a probability $s_i(t_i^\eta)$ with which the region exerts effort (chooses 1). A strategy profile $s = (s_i)$ is a vector of strategies for each region, and $s_{-i} = (s_j)_{j \neq i}$. We say that a strategy s_i is *increasing* if $s_i(t_i^\eta)$ is never decreasing in t_i^η. A strategy profile is increasing if it contains only increasing strategies.

Let $F_i^\eta(T, t^\eta; t_i)$ denote the posterior distribution function of $(T, (t_j^\eta)_{j \neq i})$ conditional on $t_i^\eta = t_i$. The expected payoff to region i, conditional on its signal t_i^ε and the vector of strategies pursued by all other regions is given by

$$u_i^\eta(x_i, s \mid t_i^\eta) = \int u_i(x_i, s(t^\eta); T) dF_i^\eta(T, t^\eta; t_i^\eta), \tag{2}$$

where with a slight abuse of notation we write $s(t^\eta)$ instead of $s_{-i}(t_{-i}^\eta)$.

Our solution concept is iterative strict dominance. We first eliminate all strategies that are strictly dominated for each region. We then then eliminate a region's strategies that are strictly dominated if the other regions are known to only play strategies that survived the initial elimination of strategies, and so on. See Osborne and Rubinstein (1994) for a textbook introduction to iterative strict dominance.

Our first main result is that uncertainty about the critical threshold T leads to selection of a unique equilibrium for the game $G(\eta)$ in which each region pursues an increasing strategy. In the limit as the noise in signals becomes arbitrarily small, the unique equilibrium strategies converge to the same strategy for all (possibly heterogeneous) regions.

Proposition 1.

(i) There is a unique strategy profile $s^* = (s_i^*)$ that survives iterated elimination of strictly dominated strategies in the game G^η. The profile $s^* = (s_i^*)$ is increasing. In particular, there exists a unique $b_i^* \in (B_0, B_1)$ such that, for all $i \in \{1, 2, ..., N\}$:

$$s_i^*(t_i) = \begin{cases} 1 & \text{if} \quad t_i \leq t_i^*, \\ 0 & \text{if} \quad t_i > t_i^*. \end{cases} \tag{3}$$

The strategy profile s^* is the unique Bayesian Nash equilibrium of G^η.

(ii) In the limit as $\eta \to 0$, all t_i^* converge to a common limit t^* given by

$$t^* = \sum_{i=1}^{N} \frac{\omega_i c_i}{b_i}. \tag{4}$$

The multiplicity of equilibria documented in Observation 1 and the literature obtains only in the extreme case of common knowledge about the entire coordination game. Even a vanishing amount of uncertainty leads to selection of a unique equilibrium. This uniqueness allows us to study the strategic behavior of regions without the ambiguity that multiple equilibria create. Using Proposition 1, we can derive a number of intuitive comparative statics that link a disease's characteristics to equilibrium behavior.

Proposition 2. *Consider the game G^η and let $\eta \to 0$.*

(i) The threshold t^* is monotone increasing in c_i, for all i;

(ii) The threshold t^* is monotone decreasing in b_i, for all i.

A higher cost of effort [benefit from control] for any region decreases [increases] the probability that a disease gets controlled in all regions. In equilibrium, a disease will be controlled if and only if $T > t^*$. Since t^* is increasing in c_i and decreasing in b_i, the ex ante (before T is drawn) probability that a disease will be controlled is decreasing in c_i and increasing in b_i, for all i.

Why should the cost-benefit ration c_i/b_i for region i affect the decision problem of another region $j \neq i$? It is intuitively clear that an increase in c_i [decrease in b_i] affects the decision of region i. Suppose that, for given t_i^η and c_i/b_i, some region i is just indifferent between exerting effort and doing nothing. Then an increase of the cost-benefit ratio c_i/b_i, given t_i^η and the strategies pursued by all other regions, will tip i's preference over from indifference to not exerting effort. This affects the other regions because exerting effort is weakly more attractive for them when region i participates compared to when it does not. Hence, regions $j \neq i$ too will become less inclined to control the disease, making them less willing to try and control the disease, too.

The comparative statics in Proposition 2 may sound obvious, but we haste to add that these results cannot be derived in a game with multiple equilibria. Indeed, as we discussed before, in a game with multiple equilibria only additional and ad hoc requirements, imposed by the modeler, can preempt the situation in which a disease with high cost-benefit ratios c_i/b_i for all i is controlled whereas one with low cost-benefit rations c_i'/b_i' is not from being an equilibrium outcome. Our global game of disease control instead offers a convincing game theoretic argument, grounded in rational behavior, to prevent such counter-intuitive situations from occurring.

Propositions 1 and 2 discuss equilibrium strategies. For practical purposes, one may be more interested in outcomes. With a slight abuse of wording, define a *pandemic* as the situation in which there is an disease in every region; hence, for a given disease (b, c, T) a pandemic occurs if and only if $X < T$. The equilibrium uniqueness established in Proposition 1 allows us to identify precise conditions under which a pandemic arises. Proposition 3 discusses these conditions and provides details about the (possible) inefficiency of a pandemic.

Proposition 3. *Consider the game G^η and let $\eta \to 0$.*

(i) For all $T > t^*$, there will *be a pandemic;*

(ii) For all $T < t^*$, there will not *be a pandemic;*

(iii) A pandemic is always inefficient.

Interestingly, though all regions favor the outcome in which a pandemic is avoided to one in which it is not, there is an entire range of diseases (all (b, c, T) such that $T < t^*$) for which rational regions will choose not to exert effort. A pandemic can hence be game theoretically rational; not because the cost of effort is too high to warrant controlling a disease (we ruled this out by assumption), but because the strategic uncertainty about other regions' efforts makes the individual gamble to control too risky in light of the coordination required. The presumption that rational regions will coordinate on the preferred equilibrium of the game of perfect information does not survive the addition of even vanishing uncertainty about payoffs.

The next corollary follows immediately from Propositions 2 and 3.

Corollary 1. *Consider the game G^η and let $\eta \to 0$. Ceteris paribus, a pandemic is more likely if i) the disease is harder to control (T higher), ii) the benefit from successful control is lower (b lower), and/or iii) the cost of effort toward control is higher (c higher).*

If we interpret b_i as a measure of how harmful a disease is, then our results suggest that more harmful diseases are more likely to be controlled and, hence, cause fewer casualties in equilibrium. This makes sense. Chickenpox, for example, is clearly pandemic but generally so mild that no efforts toward eradication are currently undertaken (McKendrick, 1995). Ebola, in contrast, is mostly restricted to West-Africa and was not allowed to spread further, partly because it is such a damaging disease (Flückiger and Ludwig, 2019). The reason for this discrepancy is that a more harmful or dangerous disease tend to warrant more caution. Similar cooling effects on behavior can be observed in the reduction in crashes on a road after speed bumps are introduced (Ewing, 1999; De Borger and Proost, 2013).

How to avoid a rational pandemic?

4 P-Delegation

The unique equilibrium identified in Proposition 1 is the only strategy vector that survives iterated elimination of strictly dominated strategies. Being a dominance solvable game, a simple pledge or promise by regions to always coordinate on the efficient outcome of G^η (avoid a pandemic) cannot work: the promised strategy to always exert effort is dominated by s_i^* for each region. All regions hence have an incentive to deviate from their promise, and this is common knowledge, so a pledge to always exert effort is incredible. Simple promises cannot prevent pandemics.

We discuss one potential solution to overcome this problem: the possibility for regions to delegate their decision making to a third agent who takes (binding) decisions on the region's behalf. We will show that classic delegation *can* work to help regions solve their coordination problem, but does not have to, as the decision to delegate is itself a coordination problem. We then introduce the novel concept of P-delegation, or delegation conditional on sufficient participation, and derive a more positive result.

Suppose there is a higher Authority (e.g. a federal government) to which the regions (e.g. individual states) can delegate their decision to exert effort in the face of the outbreak of a disease. We take a somewhat extreme position and assume that the Authority will always force the regions that delegated to exert effort. It would be reasonable to assume that the Authority has some objective function which it aims to maximize and that this maximization behavior does not *always* lead to coordinated effort by the delegating regions. However, we will show that even if the Authority rules in this extreme and, for individual regions, seemingly unfavorable way, delegation by all regions – and, consequently, successful control of any disease – can be made a dominant strategy of the P-delegation game. It follows that a more sensible Authority will also attract full coordination on P-delegation. We furthermore assume that the Authority observes (b, c, T) directly; as will become clear in the argument, it matters only for details of the analysis whether the authority observes the true T or instead receives a noisy signal of it so we simplify notation and assume that the authority observes T without error.

Formally, the problem is modeled as a two-stage game that we denote Γ^η. In stage 1 of Γ^η, all regions simultaneously decide whether to delegate their decision-making powers to the Authority. We let $\mathcal{D} \subseteq N$ denote the subset of regions that delegated their decision making powers to the Authority. After stage 1, a disease breaks out, T is drawn, and stage 2 of Γ^η begins. In stage 2, the non-committed regions simultaneously determine their effort choices in the global game $G^\eta(\mathcal{D})$ which is the global game G^η with the actions sets for all $i \in \mathcal{D}$ reduced to $\{1\}$ (by virtue of their decision to delegate). Given this setup, it is easy to see that $G^\eta(\mathcal{D})$ is simply a normalized version of the global game G^η for the regions who did *not* delegate in stage 1, i.e. all $i \in N \setminus \mathcal{D}$. We refer to the proof of Proposition 1 to establish that in $G^\eta(\mathcal{D})$, letting $\eta \to 0$, there will be a unique and common threshold signal $t^{**}(\mathcal{D})$ such that each $i \in N \setminus \mathcal{D}$ chooses $x_i = 1$ if $t_i^\eta \leq t^{**}(\mathcal{D})$ and $x_i = 0$ otherwise. If no region delegates in stage 1, stage 2 of Γ^η reduces to the global game analyzed in the previous section and $t^{**}(\emptyset) = t^*$.

We first study standard delegation, i.e. delegation that is not conditional on the realization of $\mathcal{D}$. Proposition 4 shows that the possibility to delegate decision-making powers does not solve the regions' coordination problem. Indeed, in some sense delegation further complicates it.

Proposition 4. *Let $\eta \to 0$. Both $\mathcal{D} = \emptyset$ and $\mathcal{D} = N$ can be consistent with a Perfect Bayesian equilibrium of Γ^η.*

Proof. It is not profitable to deviate when $\mathcal{D} = N$ for any region i as the region will play $x_i = 1$ both when $\mathcal{D} = N$ and in the alternative situation in which $\mathcal{D} = N \setminus \{i\}$ and i is left to choose. $\mathcal{D} = N$ is hence a Perfect Bayesian Equilibrium (PBE) of the game.

Next consider the case in which no region delegates, $\mathcal{D} = \emptyset$. Straightforward algebra yields that no region i has an incentive to unilaterally deviate from $\mathcal{D} = \emptyset$ in stage 1 (given the anticipated behavior in stage 2) if and only if

$$\frac{(1 - \omega_i)t^* - t^{**}(\{i\})}{(1 - \omega_i)(1 - t^*)} > \frac{c_i}{b_i}, \tag{5}$$

for all $i \in N$. Hence, if (5) is satisfied then no region can gain by unilaterally deviating from the coalition structure $\mathcal{D} = \emptyset$, which is therefore a PBE. ■

The possibility to delegate decision-making power to the Authority complicates, rather than solves, the coordination problem regions face. The reason is that in this kind of setup, the decision to delegate is itself coordination problem: a region will rationally want to delegate decision-making powers if and only of sufficiently many other regions do. Hence, simple delegation does not solve our original coordination problem, it only adds a second one on top.

We propose P-delegation as a solution to the coordination problem inherent in standard delegation. As before, countries once and simultaneously decide whether to P-delegate their decision-making powers; let $\mathcal{D}^P \subseteq N$ denote the subset of P-delegating regions. Like standard delegation, P-delegation puts a region's choice of effort in the hands of the Authority in a binding way. We assume again that the Authority forces P-delegating regions to always exert effort when a disease breaks out. Importantly, however, the Authority's demand is binding if and only if the minimum participation threshold P is met in the sense that $\sum_{i \in \mathcal{D}^P} \omega_i \geq P$. We denote the two-stage game with P-delegation $\Gamma^\eta(P)$.

The participation threshold P reduces the strategic risk to P-delegation in $\Gamma^\eta(P)$, compared to standard delegation in Γ^η, since a delegating region will be bound to abide by the Authority's decision only when some minimum amount of coordination ($X \geq P$) is guaranteed to be achieved. It is easy to see that for judiciously chosen P, each region expects to gain strictly from P-delegation should the critical threshold P be met, while clearly no region loses if the pledge to delegate is null.

Let $\underline{\omega} = \min_i\{\omega_i : i \in N\}$ denote the contribution to X of the region whose effort has the lowest spillover.

Proposition 5. *Consider the $\Gamma^\eta(P)$ and let $\eta \to 0$. For all $P > 1 - \underline{\omega}$, it is a weakly dominant strategy for each region i to P-delegate in $\Gamma^\eta(P)$. In the associated Perfect Bayesian equilibrium, all diseases will be controlled.*

Proof. Let D^P_{-i} denote the set of P-delegating regions except i. For all $D^P_{-i} \neq N \setminus \{i\}$, region i is indifferent between P-delegating and not P-delegating in stage 1 of $\Gamma^\eta(P)$ since in either case the threshold $P > 1 - \underline{\omega}$ is not passed; hence, all regions will end up being free to choose their effort in stage 2 whether i P-delegates or not. If $D^P_{-i} = N \setminus \{i\}$, there are two scenarios. (a) If region i P-delegates we have $\sum_{j \in D^P} = 1 > P$, the Authority's call for all regions to exert effort in stage 2 is binding and the payoff to region i is $b_i - c_i$. (b) If region i does not P-delegate, then $\sum_{j \in D^P} = 1 - \omega_i \leq 1 - \underline{\omega} < P$ and no region is bound by the Authority's demands. Regions hence play the global game G^η in stage, the equilibrium of which is described in Proposition 1. The expected (before T is drawn) payoff to region i in this game is simply $(b_i - c_i) \cdot \Pr[T > t^*] = (b_i - c_i) \sum_{j \in N} \frac{\omega_j c_j}{b_j} < b_i - c_i$. We conclude that no region i is never worse off by P-delegating, given $P > 1 - \underline{\omega}$, but is strictly better off in case $D^P_{-i} = N \setminus \{i\}$. It follows that P-delegation is a weakly dominant strategy in stage 1 of $\Gamma^\eta(P)$ in the limit as $\eta \to 0$. In the associated equilibrium, each region P-delegates and all diseases are controlled as a result. ■

The participation threshold P that allows P-delegation to work is akin to a "trigger law", i.e. a law that becomes enforced once a key change in circumstances is achieved; in this case, that key change is the participation threshold P being passed. It seems a natural condition in the context of independent, legally separated regions. In the European Union (EU), for example, major decision can only be implemented after a unanimous vote by all member states.

5 Conclusion

This paper studies disease control in a global game. Our approach stands in contrast to the existing game theoretic literature. While "eradication games" of perfect information have multiple equilibria, the global game has a unique equilibrium. This primal distinction leads to several derivative, yet important, results. First, our model can predict when an (in)efficient pandemic occurs. Second, diseases for which the benefit of successful control is higher are more likely to be controlled (or eradicated). Third, and paradoxically, diseases that are more costly a priori end up being less costly to society, precisely because these get eradicated. Our analysis is the first to embed these well-known predictions from epidemiology models within a strategic framework.

Our results clearly demonstrate that an pandemic may occur even when this is inefficient. One way to avoid this dismal outcome is to lower eradication costs. Another possibility is to reduce the set of players, for example by having countries coordinate policies through more aggregated bodies such as the European or African Unions. Third, strong social distancing or hard border closures can contribute to contain a disease. While intuitive, extant game theoretic models do not support these conclusions.

An alternative solution is strong ex ante cooperation (rather than aggregation) between players. We support this claim in a dynamic extension of our static game, in which a subset of players forms a coalition prior to the outbreak of a disease. Members of the coalition credibly commit to exert efforts towards disease control whenever the benefit is perceived to be sufficiently high. We show that credible commitment catalyzes coordination on eradication efforts and decreases the likelihood of inefficient pandemics.

There are at least two ways to think about a coalition. First, as the outcome of intensive ex ante cooperation among players. This is perhaps most intuitively, though not exclusively, thought of as international cooperation between countries through a supra-national entity such as the World Health Organization. Second, as a reduced-form description of a sequential global game, where a subset of players determines its actions first, and only then do the remaining choose theirs.

This paper offers a new perspective on the economics of disease control. Existing studies either do not rely on game theory (Kremer, 1996; Geoffard and Philipson, 1996, 1997; Gersovitz and Hammer, 2003; Epstein et al., 2008; Fenichel et al., 2011), or make the strong assumption of perfect information (Barrett, 2003). Our global game incorporates the strategic incentives underlying the spreading of a disease while taking explicit account of uncertainty. Our results provide useful insights on how societies can prepare for future disease outbreaks, highlighting the need for increased cooperation to avoid future pandemics.

PROOF OF PROPOSITION 1

For a complete proof, please consult *this version of the paper*.

References

Acemoglu, D., Johnson, S., and Robinson, J. A. (2001). The colonial origins of comparative development: An empirical investigation. *American Economic Review*, 91(5):1369–1401.

Acemoglu, D., Johnson, S., and Robinson, J. A. (2003). Disease and development in historical perspective. *Journal of the European Economic Association*, 1(2-3):397–405.

Ahituv, A., Hotz, V. J., and Philipson, T. (1996). The responsiveness of the demand for condoms to the local prevalence of AIDS. *The Journal of Human Resources*, 31(4):869–897.

Barrett, S. (2003). Global disease eradication. *Journal of the European Economic Association*, 1(2-3):591–600.

Bleakley, H. (2003). Disease and development: Evidence from the American south. *Journal of the European Economic Association*, 1(2-3):376–386.

Capasso, V. and Serio, G. (1978). A generalization of the kermack-mckendrick deterministic epidemic model. *Mathematical Biosciences*, 42(1):43 – 61.

Carlsson, H. and van Damme, E. (1993). Global games and equilibrium selection. *Econometrica*, 61(5):989–1018.

Cervellati, M., Esposito, E., Sunde, U., and Valmori, S. (2018). Long-term exposure to malaria and violence in Africa. *Economic Policy*, 33(95):403446.

Cervellati, M., Sunde, U., and Valmori, S. (2017). Pathogens, weather shocks and civil conflicts. *The Economic Journal*, 127(607):2581–2616.

Coffey, D., Geruso, M., and Spears, D. (2017). Sanitation, disease externalities and anaemia: Evidence from Nepal. *The Economic Journal*, 128(611):1395–1432.

De Borger, B. and Proost, S. (2013). Traffic externalities in cities: the economics of speed bumps, low emission zones and city bypasses. *Journal of Urban Economics*, 76:53–70.

Epstein, J. M., Parker, J., Cummings, D., and Hammond, R. A. (2008). Coupled contagion dynamics of fear and disease: mathematical and computational explorations. *PLoS ONE*, e3955.

Ewing, R. (1999). Traffic calming state of the practice slide seminar. *Institute of Transportation Engineers. Federal Highway Administration September*.

Fenichel, E. P., Castillo-Chavez, C., Ceddia, M. G., Chowell, G., Parra, P. A. G., Hickling, G. J., Holloway, G., Horan, R., Morin, B., Perrings, C., Springborn, M., Velazquez, L., and Villalobos, C. (2011). Adaptive human behavior in epidemiological models. *Proceedings of the National Academy of Sciences*, 108(15):6306–6311.

Flückiger, M. and Ludwig, Markus andÖnder, A. S. (2019). Ebola and state legitimacy. *The Economic Journal*, 129(621):2064–2089.

Frankel, D. M., Morris, S., and Pauzner, A. (2003). Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory*, 108(1):1–44.

Gatsios, K. and Karp, L. (1991). Delegation games in customs unions. *The Review of Economic Studies*, 58(2):391–397.

Geoffard, P. Y. and Philipson, T. (1996). Rational epidemics and their public control. *International Economic Review*, 37(3):603–624.

Geoffard, P. Y. and Philipson, T. (1997). Disease eradication: Private versus public vaccination. *The American Economic Review*, 87(1):222–230.

Gersovitz, M. and Hammer, J. S. (2003). The economical control of infectious diseases. *The Economic Journal*, 114(492):1–27.

Habla, W. and Winkler, R. (2018). Strategic delegation and international permit markets: Why linking may fail. *Journal of environmental economics and management*, 92:244–250.

Kremer, M. (1996). Integrating behavioral choice into epidemiological models of AIDS. *The Quarterly Journal of Economics*, 111(2):549–573.

McKendrick, M. (1995). Acyclovir for childhood chickenpox. Cost is unjustified. *The British Medical Journal*, 310(6972):108109.

Osborne, M. J. and Rubinstein, A. (1994). *A course in game theory*. MIT press.

Philipson, T. (1996). Private vaccination and public health: An empirical examination for U.S. measles. *The Journal of Human Resources*, 31(3):611–630.

Rappuoli, R. (2004). From Pasteur to genomics: progress and challenges in infectious diseases. *Nature Medicine*, 10:11771185.

Shastry, G. K. and Weil, D. N. (2003). How much of cross-country income variation is explained by health? *Journal of the European Economic Association*, 1(2-3):387–396.

Siqueira, K. and Sandler, T. (2007). Terrorist backlash, terrorism mitigation, and policy delegation. *Journal of Public Economics*, 91(9):1800–1815.

Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, 80(1):234–248.