

Selfish Drug Allocation for Containing an International Influenza Pandemic at the Onset

Peng Sun, Liu Yang

Fuqua School of Business, Duke University, Durham, North Carolina 27708
{peng.sun@duke.edu, liu.yang@duke.edu}

Francis de Véricourt

European School of Management and Technology, Berlin, Germany, devericourt@esmt.org

Recent epidemiologic studies have suggested that the prophylactic use of antiviral drugs could slow down the spread of an influenza epidemic. Because drug stockpiles are presently scattered in different countries, the outbreak of an epidemic gives rise to a game in which each country must make decisions about how best to allocate its own stockpile in order to protect its population. We develop a two-period multivariate Reed-Frost model to represent the spread of the epidemic within and across countries at its onset. We consider the first two periods only to mimic the exponential growth of an epidemic in its early stage, while keeping the model tractable. Preliminary numerical studies suggest that insights from the two-period model hold in general when considering the entire time horizon. Our model captures three critical sources of uncertainty: the number of initial infections, the spread of the disease, and drug efficacy. We show that for small probabilities of between-country infections, the underlying game is supermodular, Nash equilibrium exists, and there is a unique one that is Pareto optimal among all existing equilibria. Further, we identify sufficient conditions under which the optimal solution of a central planner (such as the World Health Organization) constitutes a Pareto improvement over the decentralized equilibrium, suggesting that countries should agree on an allocation scheme that would benefit everyone. By contrast, when the central planner's solution does not constitute a Pareto improvement, minimizing the total number of infected persons globally requires some countries to sacrifice part of their own population, which raises intriguing ethical issues.

Subject classifications: Reed-Frost model; epidemic control; influenza; supermodular games.

Area of review: Decision Analysis.

History: Received June 2008; revisions received January 2009, April 2009, May 2009; accepted May 2009.

1. Introduction

Most developed countries' preparedness plans for a possible pandemic influenza require stockpiling antiviral drugs (Rosenthal 2005). Although the efficacy of these drugs (such as Tamiflu, see Ward et al. 2005, and references therein) is not well known, recent studies have suggested that mass prophylaxis distribution of them might be able to contain a pandemic (Ferguson et al. 2005). Even if a total extinction of the epidemic might be difficult to achieve, antivirals can buy time for a vaccine to be developed, for more drugs to be produced, or for other efficient policies to be implemented. Extant imperfect vaccines (i.e., vaccines that are poor matches to virus strains that are virulent against humans) have also been shown to slow the spread of the disease (Longini et al. 2005).

The use of antivirals (or imperfect vaccines) to contain an epidemic necessitates massive administration of drugs over a very short time period, right at the onset of the first infections. However, at present drugs or vaccines are purchased mainly by the least vulnerable countries (for instance, most of world's flu vaccines are produced in Western Europe, North America, and Japan, Anonymous

2005). Regions where transmissions of the H5N1 influenza A virus¹ to humans have so far taken place, such as Southeast Asia, do not have sufficient funds to maintain large stocks of antivirals. This implies that drugs might need to be diverted to other countries if an epidemic were to be contained (Monto 2005). (In our settings, antiviral drugs and vaccines have similar effects, and we use the term "drug" to designate both of them without distinction.)

The main objective of this study is to examine how countries possessing stocks of drugs would react if an epidemic started in a country possessing little or no antiviral medication. How much of its limited supply of a drug would a country give up to contain the pandemic? How much of the drug would each country keep to guard against possible future infections? We model this problem as a game involving the countries that possess antiviral or vaccine stockpiles by developing a mathematical model that represents the dynamics of an epidemic in its early stage in a simple, but tractable way. This model captures three important sources of uncertainty: (1) the number of initial infections; (2) the spread of the disease within and between countries; and (3) the drug efficacy. We derive insights into the success (or lack thereof) at containing an outbreak when countries

hoard these drugs selfishly, that is, when they seek to protect their respective populations, even, if necessary, at the expense of others.

Our paper follows a multidisciplinary approach as it uses a game-theoretic framework with stochastic models to address the issue of containing a pandemic. To the best of our knowledge, this study constitutes the first analysis of epidemic control where antiviral drugs or vaccines are distributed by strategic agents who seek to protect different parts of the population. By contrast, previous analyses have typically considered a central planner seeking to contain an epidemic in the whole population. This line of research usually focuses on cost-benefit analysis of public health policies (Diekmann and Heesterbeek 2000, Brandeau 2004). In particular, recent works have explored the optimal allocation of limited resources in a nonhomogeneous population (see, for instance, Brandeau et al. 2003, and references therein). However, we show in this paper that ignoring strategic behaviors can lead to overoptimistic assessments of the impact of the control. That is, when decision makers are strategic, their allocation decisions might be drastically different from the centralized solution, which leads to more infections in the world compared to the centralized solution.

Recent analyses of influenza containment efforts are based mainly on extensive and detailed numerical studies. In particular, Ferguson et al. (2005) have run simulations that show how targeted mass prophylactic use of antiviral drugs significantly increases the extinction probability of an epidemic. Longini et al. (2005) and Germann et al. (2006) further suggest that the use of inefficient vaccines instead of (or in conjunction with) antivirals would also drastically slow the spread of the disease.²

For our purpose, we represent the epidemic at an aggregate level with a two-period multidimensional Reed-Frost model. The classical Reed-Frost model assumes that susceptibles (healthy individuals who could be infected) and infectives (infected individuals who could infect susceptibles) mix homogeneously within a closed population. Thus, the number of susceptibles at the end of period $t + 1$ follows a binomial distribution with probability $(1 - a)^{Y_t}$ and sequence size X_t , where a is the within-population infection probability and Y_t and X_t are the numbers of infectives and susceptibles at time period t , respectively. The infection probability a is equal to the probability that two individuals meet during the virus' latent period multiplied by the probability of infection per contact. Note that considering a stochastic model makes the analysis more challenging. However, the future spread of the disease constitutes a very important source of uncertainty when we try to contain a pandemic at the onset. A stochastic framework therefore appears necessary in our settings.

The main simplification of this model lies in the homogeneous mixing assumption, an assumption that is made in most of the epidemiology models, including deterministic SIR models (Daley and Gani 2001). Nevertheless,

the Reed-Frost model has been shown to be successful at describing epidemic dynamics (see, for instance, Daley and Gani 2001). Furthermore, homogeneous mixing is a key assumption of most tractable models of epidemics. In particular, to account for inhomogeneity, so-called spatial models (see, for instance, Ball 1991) consider multiple locations and migration of infectives, while assuming homogeneous mixing of the population in each location. The multidimensional Reed-Frost model developed in this paper may be considered as a discrete-time approximation to the continuous-time spatial model, in which the migrations and infections occur in the same time period.

We partition the world of interest into $m + 1$ countries characterized by homogeneous mixing both within and between populations that have different infection probabilities. The epidemic starts in country 0, and the remaining m countries possess drug stockpiles. In addition, we consider imperfect drugs such that the infection probability decreases when the susceptible, the infective, or both receive treatment before interacting. Thus, $2(m + 1)$ random variables are needed to represent the whole system, in which the number of susceptibles at the end of each period follows multivariate binomial distributions.

We also need to specify the objectives of each country in our model. For centralized resource allocation decisions, the cost-benefit analyses in the epidemiology literature often regard the goal as being to decrease the reproduction number, R_0 (for a stratified population, defined as the spectral radius of the so-called "next generation matrix," see, for example, Hill and Longini 2003), to below 1 (Longini et al. 2005). In a decentralized world like the one that we model in this paper, it is not entirely clear what an appropriate condition for R_0 would be for each individual country, because the condition should allow for an outbreak in the world outside the considered country. Therefore, in this paper we take a different approach by considering the main objective for mass drug administration to be either the eradication of the disease or to buy time at the onset of the epidemic (as argued by Mills et al. 2006). In our case, the early stage of the pandemic should occur in the first few periods of our Reed-Frost model. We consider the first two periods only, which allows one to mimic the exponential growth of the epidemic in its early stage, while keeping the model tractable. Indeed, when considering the first period only, all countries keep their drugs for their own populations, because potential benefits have not materialized during a single time period. By contrast, numerical studies suggest that our findings for two periods hold in general when considering the entire time horizon of the epidemic, with the difference that in this case countries are willing to give up their drugs more often. In other words, the two-period model appears to predict that countries act more selfishly than they would if they were to consider the final size of their population.

When the objective is to limit the impact of the pandemic, we assume that each country seeks to minimize the

expected number of infections at the end of the two periods. Because population sizes are fixed in our model, this objective is equivalent to maximizing the number of susceptibles (i.e., noninfected persons) in the two first periods. Besides the expected value objective, we also study another plausible objective for each country—minimizing the probability of the disease spreading over from other countries in the beginning (two first periods) of the epidemic. This is also equivalent to maximizing the probability that no individual within the borders is infected during the first two periods. Other objective functions may exist that estimate the total economic costs of a pandemic and its impact on the quality-adjusted life years (see, for instance, Brandeau et al. 2003, for a study on the centralized resource allocation problem). In case of an outbreak, however, we believe that the number of infected individuals and the probability of the disease spreading over from other countries are the most pressing concerns of the decision makers.

At the beginning of the first time period, an infective is discovered in country 0. All countries share a common subjective probability distribution for the number of infected persons present in the population. In particular, this distribution should take into account the possibility of multiple introductions of pandemic-capable strains, as discussed by Mills et al. (2006). Based on this distribution and on the different parameters of the system (country sizes, infection probabilities, and inventory levels), each country makes its drug allocation decisions. A given country (say, country i) prefers the rest of the world to contribute more drugs to country 0, so that it could keep its drug stockpile to protect its own population. On the other hand, if country 0 does not receive enough drugs, country i may choose to give up part of its drug stockpile to help contain the pandemic in other populations. Indeed, if the epidemic is not contained in the country where it first occurs, it might quickly spread to the rest of the world, dramatically increasing the chance of infection from abroad. This situation thus gives rise to a game in which each country seeks to protect its population by strategically allocating its drugs.

Several papers have considered rational individual decision making in the face of epidemics (see Philipson 2004, and references therein). Bauch and Earn (2004) have recently introduced game theory into epidemic models. Their focus, however, is on voluntary vaccination, which differs from the topic of this paper. They described a “commons dilemma” such that the risk associated with receiving vaccination is certain and individual specific, whereas the return is both uncertain (i.e., the individual may never come into contact with an infective) and widely distributed. Thus, rational self-interest might incline an individual to avoid vaccination, which, in turn, could prevent complete eradication of the disease. Chick et al. (2008) also developed a game theory framework to analyze the coordination of vaccine production where actors (the government and manufacturers) have different incentives. Their work contributes to the line of research concerned with the production and

stockpiling of vaccines (see also Deo and Corbett 2009). More generally, the challenge of containing influenza outbreaks has received a lot of attention recently in the OR/MS literature (see Wu et al. 2005, Kornish and Keeney 2008, Cho 2009, and references therein).

Our model also links to the economics literature concerning the voluntary contributions to public goods. This area of research is primarily interested in understanding how different agents may try to enjoy a “free ride,” i.e., benefit themselves at the expense of others, while encouraging the latter to increase their own contributions to some public project or good. For instance, Marx and Matthews (2000) provide sufficient conditions under which the final public project completion time reduces to zero, and free riding vanishes. In our settings, the “public project” could be the worldwide eradication of the epidemic. The “contributions” would then correspond to the number of drugs allocated to the country where the first infective appeared. However, situations occur where all players are better off not achieving the public good (i.e., not giving up drugs to country 0). In these cases, the centralized solution basically requires countries to “sacrifice” part of their population in order to save more individuals in country 0. Protecting this population is therefore not quite equivalent to a public good. In our case, the inefficiency of a decentralized solution is not necessarily due to free riding, but rather due to the misalignment of objectives of players and that of the central planner.

Furthermore, when all contribution decisions are made at the beginning of the time horizon and for linear payoff functions, the problem of free riding reduces to a prisoners’ dilemma and no agent contributes to the public project. However, this is not the case when the project corresponds to containing a global pandemic as shown in this paper. Actually, we identify situations where countries are willing to give up all their drugs to cover the part of the population where the first infectives appeared. This significant difference is due to the dynamics of an epidemic spread that induces nonlinear payoff functions that are neither concave nor convex. In particular, we show that countries either keep or give away all their drugs (and do not split their resources).

For the two-period model, we show that for small between-country infection rates, the underlying game is supermodular, which implies the existence of a unique Nash equilibrium that is Pareto optimal among all existing equilibria. This equilibrium is characterized by a subset of countries that give up all their drugs to contain the spread of the epidemic in country 0, whereas all other countries keep their stockpiles to protect their own populations. If a central planner such as the World Health Organization (WHO) were to make the allocation decision in order to minimize the total average number of infectives in the world (irrespective of nationalities), all existing stockpiles would be used to supply country 0. We identify sufficient conditions under which the optimal solution of the central planner

constitutes a Pareto improvement over the decentralized equilibrium. In particular, for large populations and a high probability of disease transmission in country 0, using all existing drugs to contain the epidemic in country 0 would benefit all countries. On the other hand, when countries act selfishly and seek to protect their own populations, there could be human cost in the rest of the world. These findings raise important ethical issues, especially when the centralized decision does not constitute a Pareto improvement over the decentralized situation. In this latter case, minimizing the total number of infectives in the world requires some countries to sacrifice part of their own population. Our numerical studies show that these insights hold, in general, in the entire horizon case.

We present our epidemic model and the drug allocation game in §2. The case where countries maximize the expected final number of susceptible in their population is analyzed in §3. The comparison with the decisions of a central planner follows in §3.3. Section 4 extends some of our results to the case where countries maximize the probability of not being infected from abroad. Section 5 numerically tests some of the model assumptions. In particular, we explore how our main insights extend when considering the entire time horizon of the epidemic. We draw several conclusions and point to future research directions in §6.

2. The Drug Allocation Game

Consider $m + 1$ countries, $0, 1, \dots, m$ where country i 's population size is equal to N_i for $0 < i \leq m$ and $N_0 + 1$ for country $i = 0$. The transmission probability from a nontreated infective in country i to a nontreated susceptible in country j is denoted by $a_{ij} \geq 0$. Access to drugs is assumed to decrease susceptibility by a factor δ , and infectiousness if infected by a factor ξ . For instance, recent statistical studies estimate that antiviral drugs reduce susceptibility to infection for an uninfected individual by 30%, and infectiousness for infected individuals by 60%, which corresponds to $\delta = 70\%$ and $\xi = 40\%$ (Ferguson et al. 2003, Yang et al. 2006). When both parties have used drugs, the probability of infection per contact decreases by $\sigma := \delta\xi$. In what follows, we refer to a_{ij} as the probability of between-country infection when $i \neq j$ and as the probability of within-country infection when $i = j$. We assume that the disease spreads more easily within a country than between countries; that is, $a_{ij} < a_{ii}$ for all $i \neq j$.

We consider a multidimensional Reed-Frost epidemic model of two periods. Let X_i^t represent the number of susceptibles in location i at time t who have not been treated, whereas \bar{X}_i^t represents the number of susceptibles at time t who have taken the drug previously. Similarly, \bar{Y}_i^t and Y_i^t represent the number of infectives in location i at time t who have (\bar{Y}_i^t) and have not (Y_i^t) used the drug, respectively. Thus, $Y_i^t + \bar{Y}_i^t$ individuals spread the disease during period t before becoming symptomatic and are removed

by the end of the period. The number of susceptibles and infectives therefore satisfy the following relationships,

$$X_i^t = X_i^{t-1} - Y_i^t, \quad \bar{X}_i^t = \bar{X}_i^{t-1} - \bar{Y}_i^t. \quad (1)$$

Following the Reed-Frost model assumptions, the numbers of susceptibles have binomial distributions. In a given period, susceptibles of country i can be infected by infectives of the same country and of another country j . We can then easily deduce the following conditional probabilities,

$$(X_i^{t+1} | X_i^t, \{Y_j^t, \bar{Y}_j^t\}_j) \sim \mathcal{B}\left(X_i^t, (1 - a_{ii})^{Y_i^t} (1 - \xi a_{ij})^{\bar{Y}_i^t} \cdot \prod_{j \neq i} (1 - a_{ji})^{Y_j^t} (1 - \xi a_{ji})^{\bar{Y}_j^t}\right), \quad (2)$$

$$(\bar{X}_i^{t+1} | \bar{X}_i^t, \{Y_j^t, \bar{Y}_j^t\}_j) \sim \mathcal{B}\left(\bar{X}_i^t, (1 - \delta a_{ii})^{Y_i^t} (1 - \sigma a_{ii})^{\bar{Y}_i^t} \cdot \prod_{j \neq i} (1 - \delta a_{ji})^{Y_j^t} (1 - \sigma a_{ji})^{\bar{Y}_j^t}\right), \quad (3)$$

where $\mathcal{B}(N, p)$ represent a Binomial distribution with parameters N and p .

Because of possible multiple introductions or early infections, several infected individuals might already exist in the country at the time the first case is discovered. We assume that countries share a common subjective distribution for Y , the number of infected individuals in country 0 at time $t = 0$. This implies that the multiple introductions and first infections take place in country 0 only. Thus, $Y_0^0 = Y$, $Y_i^0 = 0$ for $i > 0$, and $\bar{Y}_i^0 = 0$ for all i . The initial conditional probabilities are then equal to $(X_i^1 | X_i^0, Y) \sim \mathcal{B}(X_i^0, (1 - a_{0i})^Y)$ and $(\bar{X}_i^1 | \bar{X}_i^0, Y) \sim \mathcal{B}(\bar{X}_i^0, (1 - \delta a_{0i})^Y)$ for $i = 0, 1, \dots, m$. We further let $E_Y(\cdot)$ and $\varphi_Y(\cdot)$ be the expectation operator and characteristic function, respectively, of the random variable Y .

Each country $i > 0$ has a stockpile of drugs that can treat up to K_i susceptibles. For simplicity of exposition, we assume country 0 does not have any drug stockpile, or $K_0 = 0$. Our results still hold for $K_0 > 0$. We assume that $K_i \leq N_i$, and the world's total drug supply $K := \sum_j K_j$ is less than the population of country 0, that is, $K \leq N_0$. This assumption appears to hold under the current situation, where a pandemic flu is likely originating from one of the populous Southeast Asian countries. Note also that similar constraints on the number of drugs are often placed in simulation of epidemics (see, for example, Ferguson et al. 2005). Technically, such an assumption helps to alleviate some of the difficulties in showing the supermodularity of the game.

We denote by n_j^i the number of country j 's susceptibles, $j \geq 0$, treated by drugs from country i 's stockpile (when drugs represent antivirals and not vaccines, a unit of inventory represents the number of doses required to treat an individual during the two periods). For simplicity, we assume that the unidentified first infectives Y are not treated when drugs are distributed (similar results hold when they

all receive treatment). Given all allocation decisions $\{n_j^i\}_{i,j}$ with $\sum_j n_j^i = K_i$, the initial numbers of susceptibles are equal to

$$X_i^0 = N_i - \sum_{k=1}^m n_i^k \quad \text{and} \quad \bar{X}_i^0 = \sum_{k=1}^m n_i^k, \quad \text{for } i > 0, \quad \text{and} \quad (4)$$

$$X_0^0 = N_0 - Y - \sum_{k=1}^m n_0^k \quad \text{and} \quad \bar{X}_0^0 = \sum_{k=1}^m n_0^k.$$

In the following we also use \mathbf{n} to denote the decision matrix $\{n_j^i\}_{(m+1) \times (m+1)}$ such that row vector \mathbf{n}^i represents country i 's decisions.

3. Expected Value Case

In this section, we study the case in which each country is risk neutral and seeks to minimize the average number of total infectives during the two periods. Because in our model we assume a fixed population size, the objective is equivalent to maximizing the average number of noninfected susceptibles at the end of the time horizon. In §4, we extend our results to the case where the objective function is the probability of no infection, which might describe more accurately some countries' actual motivations.

The problem, however, is generally intractable. A country's response curves are neither convex nor concave and sometimes possess multiple local maxima. On the other hand, we are primarily interested in contaminations between countries at the onset of the epidemic for which the between-country infection rates $\{a_{ij}\}_{i \neq j}$ are small. It turns out that in this case the problem is well structured. The key step in our analysis consists of taking Taylor expansions (in the between-infection rates) of the country's average number of susceptibles in the end of period 2. Our numerical results (see Online Appendix A) suggest that the structure of the game holds even when the between-country infection rates are not too small. Online Appendix A is available as part of the online version that can be found at <http://or.journal.informs.org/>.

3.1. Country i 's Response Curve

We next derive country i 's response curve. To that end, we first evaluate the average number of susceptibles in the end of period 2 given the initial treated and untreated population sizes for country i . To simplify our mathematical expressions, we introduce the following quantities, $A_{ji} = a_{ji}[1 - (1 - a_{0j})^Y]$ and $B_{ji} = a_{ji}[1 - (1 - \delta a_{0j})^Y]$. Note that A_{ji} and B_{ji} depend on Y , but for the sake of simplicity we do not make this dependence explicit in our notation.

PROPOSITION 1. For country $i = 0, 1, \dots, m$,

$$E[X_i^2 \mid \{X_j^0, \bar{X}_j^0\}_{j \geq 0}, Y]$$

$$= X_i^0 \frac{(1 - a_{0i})^Y}{1 - A_{ii}} \prod_{j=0}^m (1 - A_{ji})^{X_j^0} (1 - \xi B_{ji})^{\bar{X}_j^0},$$

$$E[\bar{X}_i^2 \mid \{X_j^0, \bar{X}_j^0\}_{j \geq 0}, Y]$$

$$= \bar{X}_i^0 \frac{(1 - \delta a_{0i})^Y}{1 - \sigma B_{ii}} \prod_{j=0}^m (1 - \delta A_{ji})^{X_j^0} (1 - \sigma B_{ji})^{\bar{X}_j^0}.$$

PROOF. See Online Appendix B. \square

Country i seeks to maximize the average number of susceptibles in its population at the time horizon that corresponds to the following optimization problem,

$$\max_{\mathbf{n}} f_i(\mathbf{n}) := E_Y[E[X_i^2 + \bar{X}_i^2 \mid \{X_j^0, \bar{X}_j^0\}_{j \geq 0}, Y]],$$

subject to the following constraints

$$\sum_j n_j^i = K_i, \quad 0 \leq n_j^i \leq N_j - \sum_{k \neq i} n_j^k.$$

Following Equation (4) and Proposition 1, we obtain

$$f_i(\mathbf{n}) = \left(N_i - \sum_{k=1}^m n_i^k \right) E_Y[G_1^i] + n_i^0 \sum_{k=1}^m n_i^k E_Y[G_2^i], \quad (5)$$

where

$$G_1^i = \frac{(1 - a_{0i})^Y}{1 - A_{ii}} (1 - A_{0i})^{-Y}$$

$$\cdot \prod_{j=0}^m (1 - A_{ji})^{N_j - \sum_{k=1}^m n_j^k} (1 - \xi B_{ji})^{\sum_{k=1}^m n_j^k},$$

$$G_2^i = \frac{(1 - \delta a_{0i})^Y}{1 - \sigma B_{ii}} (1 - \delta A_{0i})^{-Y}$$

$$\cdot \prod_{j=0}^m (1 - \delta A_{ji})^{N_j - \sum_{k=1}^m n_j^k} (1 - \sigma B_{ji})^{\sum_{k=1}^m n_j^k}.$$

$E_Y[G_1^i]$ and $E_Y[G_2^i]$ represent the proportion of susceptibles in the population who were and were not treated, respectively.

We are now ready to study country i 's optimal drug allocation given the decisions made by all other countries. As will be seen, most of our proofs are based on treating the $f_i(\mathbf{n})$ as continuous functions. Most of our results also rely on the approximation that the between-country transmission probabilities a_{kl} , $k \neq l$, are "small enough" or "sufficiently small." Specifically, it means that there exists an $\epsilon > 0$ such that if all $a_{kl} < \epsilon$, then the corresponding result holds. Note that the transmission probability a_{kl} reflects the probability of contact and infection between any infective in country k and any susceptible in country l . In practice a_{kl} should be very small.

Our first result states that country i contributes drugs to country 0 only.

PROPOSITION 2. For small-enough between-country transmission rates a_{kl} , $k \neq l$, country i never gives drugs to another country $j \neq 0$.

PROOF. See Online Appendix C. \square

The proof (Online Appendix C) reveals the intuition behind the result: From country i 's perspective, the marginal benefit from giving drugs to country 0 always exceeds the marginal benefit from giving drugs to another

country $j \neq 0$. Therefore, the question is whether to use the drug within country i or in country 0.

Denote, then, n_i to be the number of drug units that country i decides to keep for itself, so that $K_i - n_i$ is allocated to country 0. It follows that, with $K := \sum_j K_j$,

$$X_i^0 = N_i - n_i, \quad \bar{X}_i^0 = n_i,$$

$$X_0^0 = N_0 - Y - K + \sum_{k=1}^m n_k, \quad \bar{X}_0^0 = K - \sum_{k=1}^m n_k.$$

With drugs from country i only going to countries i and 0, country i 's objective function becomes

$$f_i(\mathbf{n}) = (N_i - n_i)E_Y[G_1^i] + n_i E_Y[G_2^i] \\ = N_i E_Y[G_1^i] + n_i (E_Y[G_2^i] - E_Y[G_1^i]), \quad (6)$$

where

$$G_1^i = \frac{(1 - a_{0i})^Y}{1 - A_{ii}} (1 - A_{0i})^{N_0 - Y - K + \sum_j n_j} (1 - \xi B_{0i})^{K - \sum_j n_j} \\ \cdot \prod_{j \neq 0} (1 - A_{ji})^{N_j - n_j} (1 - \xi B_{ji})^{n_j}, \quad (7)$$

$$G_2^i = \frac{(1 - \delta a_{0i})^Y}{1 - \sigma B_{ii}} (1 - \delta A_{0i})^{N_0 - Y - K + \sum_j n_j} (1 - \sigma B_{0i})^{K - \sum_j n_j} \\ \cdot \prod_{j \neq 0} (1 - \delta A_{ji})^{N_j - n_j} (1 - \sigma B_{ji})^{n_j}. \quad (8)$$

(Note that $E_Y[G_2^i]$ and $E_Y[G_1^i]$ still represent the proportion of susceptibles in the population who were and were not treated, respectively, and $f_i(\cdot)$ represents the average number of susceptibles. The mathematical expressions for $f_i(\cdot)$, G_1^i , and G_2^i are different than their counterparts in (5) due to the simplification of players' decision sets. We continue using the same notation here for the sake of clarity.)³

In the sequel, we assume the following technical condition holds for all countries i ,

$$N_i - (1 - \delta\xi)K_i \geq \frac{1 - \delta^2\xi}{1 - \delta}. \quad (9)$$

This assumption is not very restrictive. It merely rules out parameters such that δ and ξ are simultaneously very close to 1 and 0, respectively. Under reasonable choices of parameters, such as $\delta = 70\%$ and $\xi = 40\%$, the condition translates to $N_i > 0.72K_i + 2.7$, which is easily satisfied.

The optimal decision for country i given all other countries' drug allocations has a "bang-bang" structure as stated by the following result.

THEOREM 1. *For small-enough between-country transmission probabilities a_{kl} , $k \neq l$, country i 's optimal decision given all other countries' actions is always "bang-bang": either give up everything to country 0 or give nothing. Furthermore, there is a quantity a^* defined as*

$$a^* := \frac{1 - \xi - \varphi_Y[1 - a_{00}] + \xi\varphi_Y[1 - \delta a_{00}]}{(1 - \sigma)E_Y[Y]}, \quad (10)$$

such that:

1. When $a_{ii} \geq a^*$, f_i is monotonically increasing in n_i and country i keeps all its drugs for itself; and
2. When $a_{ii} < a^*$, f_i is convex in n_i .

PROOF. See Online Appendix D. \square

Quantity a^* is a common threshold for within-country transmission probability a_{ii} for all countries i . In Online Appendix E, we provide sufficient conditions for a country to be willing to give up its drugs. The "bang-bang" structure of the optimal response is rather surprising. It greatly simplifies the action profile that we need to consider when studying the equilibria of the game.

3.2. Supermodular Game and Nash Equilibrium

Theorem 1 indicates that the best strategy for country i such that $a_{ii} \geq a^*$ is degenerate with $n_i = K_i$. In other words, country i has no incentive to give up drugs to another country where the disease transmission probabilities are sufficiently low. On the other hand, when $a_{ii} < a^*$, there exist reasonable model parameters under which a country gives up all its drugs to country 0 (Online Appendix E). Hence, the drug allocation game is not degenerate and is only played by countries in which the disease transmission probability is lower than a certain threshold.

The following theorem constitutes our main result and states the existence of a unique Nash equilibrium that is Pareto optimal among all Nash equilibria.

THEOREM 2. *For small-enough between-country transmission probabilities a_{kl} , $k \neq l$, the game is supermodular among countries with $a_{ii} < a^*$, in which a^* is defined in Equation (17) of Theorem 1. Further, a unique Nash equilibrium exists that is Pareto optimal among all existing Nash equilibria. This equilibrium is such that the number of countries giving up drugs to country 0 is the highest among all existing equilibria.*

The proof depends on the supermodular structure of the game. The complete proof is in Online Appendix F. In what follows, we refer to the unique Nash equilibrium that is Pareto optimal among all existing Nash equilibria as the "Pareto-optimal Nash equilibrium."

According to Theorem 1, an equilibrium is characterized by a partition $S \cup \bar{S} = [1, \dots, m]$ such that $i \in S$ if and only if country i keeps all its drug, i.e., $n_i = K_i$. From Theorem 2, a Pareto-optimal equilibrium exists and is characterized by the set \bar{S} with the highest cardinality.

It is worth pointing out that the existence of a Pareto-optimal Nash equilibrium does not automatically guarantee it will be the outcome of the game. As pointed out in the experimental economics literature (see, for instance, VanHuyck et al. 1990, 1991), even the Pareto-dominant equilibrium that "should" be chosen according to risk dominance (Harsanyi and Selten 1988) may not occur in practice. A more recent literature on the so-called "global games" (Carlsson and Van Damme 1993) attempts to address the equilibrium selection issue by first relaxing the implicit common knowledge assumption of complete information games so that payoffs are observed with a small amount of noise. Under certain conditions, there is a unique

equilibrium in the game with noise, which serves as a prediction for the underlying complete information game that contains multiple equilibria. Frankel et al. (2003), which studies supermodular games in this framework, is probably the global games literature most relevant to our problem. However, the results in Frankel et al. (2003) do not appear to directly apply to our setting.

3.3. Central Planner

Consider now a central planner, such as the WHO, whose objective is to allocate the total world drug inventory $K = \sum_i K_i$ in order to maximize the average number of susceptibles at the end of period 2 in the world regardless of the country of residence. Let n_i be the amount of drugs the central planner allocates to country $i = 0, \dots, m$, with $n_0 = K - \sum_i n_i$, and define $\mathbf{n} = \{n_i\}_{i \geq 0}$ to be the corresponding drug allocation policy.

The central planner’s optimization problem is, therefore,

$$\max_{\mathbf{n} \geq 0, \sum \mathbf{n} = 0} f(\mathbf{n}) := \mathbb{E} \left[\sum_{i=0}^m X_i^2 + \bar{X}_i^2 \mid \{N_j - n_j, n_j\}_{j \geq 0} \right], \quad (11)$$

and from Proposition 1 we obtain,

$$f(\mathbf{n}) = \mathbb{E}_Y \left[\left(N_0 - Y - K + \sum_k n_k \right) G_1^0 + \left(K - \sum_k n_k \right) G_2^0 + \sum_{i=1}^m (N_i - n_i) G_1^i + n_i G_2^i \right], \quad (12)$$

where

$$G_1^0 = \frac{(1 - a_{00})^Y}{1 - A_{00}} (1 - A_{00})^{N_0 - Y - K + \sum_k n_k} (1 - \xi B_{00})^{K - \sum_k n_k} \cdot \prod_{j \neq 0} (1 - A_{j0})^{N_j - n_j} (1 - \xi B_{j0})^{n_j};$$

$$G_2^0 = \frac{(1 - \delta a_{00})^Y}{1 - \sigma B_{00}} (1 - \delta A_{00})^{N_0 - Y - K + \sum_k n_k} (1 - \sigma B_{00})^{K - \sum_k n_k} \cdot \prod_{j \neq 0, i} (1 - \delta A_{j0})^{N_j - n_j} (1 - \sigma B_{j0})^{n_j}.$$

When between-country infection rates are low, the central planner always supplies country 0 as much as possible.

PROPOSITION 3. *For small-enough between-country transmission rates a_{kl} , $k \neq l$, the optimal allocation policy uses all existing drugs for country 0, i.e., $n_0 = K$ and $n_i = 0$, $i > 0$.*

PROOF. See Online Appendix G. \square

Allocating all existing drugs to country 0 maximizes the number of susceptibles at the end of period 2 in the world, but might hurt some countries that could save more of their inhabitants if they acted selfishly and kept their drugs. The following results provide conditions under which supplying country 0 with all available drugs constitutes a Pareto improvement over the previous optimal decentralized equilibrium. In this case, all countries would benefit from letting the central planner decide.

PROPOSITION 4. *For $\{a_{kl}\}_{k \neq l}$ small enough, consider the decentralized equilibrium characterized by the set S of countries that do not give up their drugs (i.e., $i \in S \Leftrightarrow n_i = K_i$). For each $i \in S$, there is a population size threshold \hat{N}_i such that supplying country 0 with all existing drugs (i.e., $\sum_{j=1}^m n_j = K$) is a Pareto improvement over the decentralized equilibrium if for all $i \in S$,*

$$N_i > \hat{N}_i \quad \text{and} \quad K_i < \rho_i \left(K - \sum_{j \in \bar{S}} K_j \right),$$

where \hat{N}_i is independent of N_j , $j > 0$, and

$$\rho_i = \frac{a^*(1 - \sigma)}{a_{ii}(2 - \sigma - \delta)}.$$

PROOF. See Online Appendix H. \square

(A cumbersome analytical expression exists for the population threshold \hat{N}_i and can be found in Equation (22) in Online Appendix H.)

In other words, if countries in S are large enough and have relatively small drug inventories, giving up all existing drugs to country 0 improves the average number of noninfected susceptibles in all countries. A central planner could then try to have all countries in S agree on this policy. Without monetary enforcement, however, the centralized optimal solution may never be an equilibrium. On the other hand, the WHO may serve as a correlation center in a correlated equilibrium solution concept (Myerson 1991), which provides a coordination mechanism that potentially improves upon a Nash equilibrium in a full-information game. Unfortunately, correlated equilibria do not constitute such an improvement in our settings.⁴

The second condition $K_i < \rho(K - \sum_{j \in \bar{S}} K_j)$ is actually always satisfied if $\rho > 1$, which is reasonable when the epidemics spread sufficiently faster in country 0 than in the rest of the world. Hence, Proposition 4 implies that when all countries are large enough and the epidemic spreads sufficiently fast in country 0, allocating all existing drugs to country 0 constitutes a Pareto improvement over any decentralized equilibrium.

4. Alternative Objective: Probability of No Infection

Existing epidemiological studies (see, for instance, Ferguson et al. 2005) have focused their performance analysis on the probability of epidemic extinction in a given world area. In this section, we consider the perspective of countries with drug stockpiles that decide to maximize the probability of no infection from the country where the epidemic has started. Note that this objective is neither risk seeking nor risk averse and is not well justified from a theoretical perspective. On the other hand, maximizing the probability of no infection might more accurately represent actual responses to an outbreak, given the psychological reactions and panic that most likely would occur (see for instance Lau et al. 2006).

In order to keep the model tractable, we approximate the binomial distribution of Y_i^1 by a Poisson distribution with arrival rate $\lambda_i = a_{0i}X_i^0$. This approximation is accurate when X_i^0 is large and a_{0i} is very small. Similarly, we assume that \bar{Y}_i^1 follows a Poisson distribution with parameter $\bar{\lambda}_i = \delta a_{0i}\bar{X}_i^0$. Country i may then maximize the likelihood of no infection, which, after some algebraic simplifications, is approximately equal to

$$E_Y[P(X_i^2 = X_i^0 \cap \bar{X}_i^2 = \bar{X}_i^0 \mid \{X_j^0, \bar{X}_j^0\}_j)] \\
 \approx E_Y \left\{ (1 - a_{0i})^{YX_i^0} (1 - \delta a_{0i})^{Y\bar{X}_i^0} \cdot \exp \left\{ \sum_{j \neq i} [(1 - (1 - a_{0j})^Y) \cdot ((1 - a_{ji})^{X_i^0} + (1 - \delta a_{ji})^{\bar{X}_i^0} - 2)X_j^0 + (1 - (1 - \delta a_{0j})^Y)((1 - \xi a_{ji})^{X_i^0} + (1 - \sigma a_{ji})^{\bar{X}_i^0} - 2)\bar{X}_j^0] \right\} \right\}.$$

Country i 's optimization problem becomes, following Equation (4),

$$\max_{\mathbf{n}^i: \mathbf{n}^i \geq 0, \sum_j n_j^i = K_i} f_i(\mathbf{n}) \\
 := E_Y \left\{ (1 - a_{0i})^{Y(N_i - \mathbf{n}_i)} (1 - \delta a_{0i})^{Y\mathbf{n}_i} \cdot \exp \left\{ \sum_{j \neq i} [(N_j - \mathbf{n}_j)(1 - (1 - a_{0j})^Y) \cdot ((1 - a_{ji})^{N_i - \mathbf{n}_i} + (1 - \delta a_{ji})^{\mathbf{n}_i} - 2) + \mathbf{n}_j(1 - (1 - \delta a_{0j})^Y) \cdot ((1 - \xi a_{ji})^{N_i - \mathbf{n}_i} + (1 - \sigma a_{ji})^{\mathbf{n}_i} - 2)] \right\} \right\}, \quad (13)$$

where $\mathbf{n}_j = \sum_{k=1}^m n_j^k$ is the total number of drugs that country j receives.

It turns out that the structure of the drug allocation game is very similar to the expected value case, as shown by the following result.

THEOREM 3. *For between-country transmission probabilities a_{0k} ($k > 0$) small enough,*

1. *Country i never gives drugs to country $j \neq 0$.*
2. *Country i 's optimal decision given all other countries' actions is always "bang-bang:" either give up everything to country 0 or give nothing.*
3. *The game is supermodular; a unique Nash Equilibrium exists that is Pareto optimal among all existing equilibria, and has the highest number of countries giving drugs to country 0.*

PROOF. See Online Appendix I. \square

Note also that when N_i is large enough, country i gives all its drugs to country 0 (see Proposition 6 in Online Appendix I). Hence, although the probability of no infection constitutes a rather different objective function compared to the expected value, our main results and insights hold.

5. Numerical Study

In this section we numerically explore the present results from a computational perspective. First, in §5.1, we numerically examine the parameter settings under which having every country give up their drugs to country 0 constitutes either the Pareto-optimal Nash equilibrium or a Pareto improvement from an equilibrium. Then, §5.2 is devoted to a simulation study of the case when countries' objectives are the expected total number of infectives in the entire time horizon of the epidemic.

In the computational study presented in this section, we fix the drug effectiveness parameters $\delta = 0.7$ and $\xi = 0.4$ (Ferguson et al. 2003, Yang et al. 2006). For simplicity, we also fix the number of initial infectives $Y = 1$.

5.1. To Give or Not to Give?

In this section we numerically study whether having every country give up its drugs to country 0 constitutes either a Pareto equilibrium or a Pareto improvement from an equilibrium. To that end we conduct an extensive computational study with varied model parameters as indicated in Table 1.

Specifically, we vary the number of drug-possessing countries, m , to be 1, 4, and 7. We also vary each country's population, positing the values 10^5 , 10^6 , or 10^7 , assuming all drug-possessing countries' populations are the same. The world's total drug stockpiles, K , is assumed to be 10%, 50%, or 90% of country 0's population, and evenly distributed among all drug-possessing countries. Country 0's disease transmission probability a_{00} takes four different values, $1.3/N_0$, $2/N_0$, $2.5/N_0$, and $3/N_0$. If we consider country 0 being isolated from the rest of the world, $a_{00}N_0$ may be perceived as the basic reproductive ratio R_0 in the corresponding Reed-Frost model. The above choices of a_{00} reflect, roughly speaking, R_0 taking values 1.3, 2, 2.5, and 3. The within-country infection rates, a_{ii} , $i > 0$, are chosen to be a_{00} , $a_{00}/5$, or $a_{00}/10$. Finally, a_{ij} takes the values 1, $1/5$, $1/10$, or $1/100$ of a_{ii} .

Table 1. Model parameters.

m	1	4	7	
N_0	10^5	10^6	10^7	
N_i	10^5	10^6	10^7	
K/N_0 (%)	10	50	90	
a_{00}	$1.3/N_0$	$2/N_0$	$2.5/N_0$	$3/N_0$
a_{ii}	a_{00}	$a_{00}/5$	$a_{00}/10$	
a_{ij}	a_{ii}	$a_{ii}/5$	$a_{ii}/10$	$a_{ii}/100$

In total, we vary 7 model parameters. There are, altogether, 3,888 possible combinations of model parameters. Among the 3,888 cases, some contradict our model assumptions. For example, when $N_0 = 10^7$ and $N_i = 10^5$, the per-country stockpile, $K_i = K/m$, will always be greater than N_i . We therefore ignore such infeasible model parameters, which leaves 2,976 feasible combinations of model parameters. Of these, 1,586 (53.29%) of the cases are such that having all countries give their drugs to country 0 is the Pareto-optimal Nash equilibrium. Of the remaining 1,390 cases where the Pareto-optimal Nash equilibrium has all drug-possessing countries keep their drugs for themselves, 606 (43.60% of the 1,390 cases) are such that having all countries give their drugs to country 0 constitutes a Pareto improvement over the equilibrium. Another point worth mentioning is that in the 606 cases of Pareto improvement, the model parameters of 600 cases (99.01%), when we ignore the positive a_{ij} s, satisfy the sufficient conditions in Proposition 4. That is,

$$N_i > \left(a_{ii}(\delta\sigma - 1) + (1 - \delta) \left[1 + N_0[1 - \varphi_Y(1 - a_{00})]/E_Y[Y] - (1 - \sigma) \left(\sum_{j \in \bar{S}} K_j a^* + K_i a_{ii} \right) \right] \right) \cdot \left(a^*(1 - \sigma) \left(K - \sum_{j \in \bar{S}} K_j \right) / K_i - a_{ii}(2 - \sigma - \delta) \right)^{-1}$$

and

$$K_i < \frac{a^*(1 - \sigma)}{a_{ii}(2 - \sigma - \delta)} \left(K - \sum_{j \in \bar{S}} K_j \right).$$

In other words, our numerical study suggests that in roughly half of the cases, countries' selfish interests lead to the same drug allocation that a central planner would make. For 25% of the cases, countries playing the game will not supply country 0 with drugs, but negotiations beforehand could yield the optimal centralized decision that consists a Pareto improvement. For the remaining 25% of cases, however, country 0 will not receive drugs.

Table 2 illustrates how the model parameters affect the Pareto-optimal Nash equilibrium.

Each cell in Table 2 indicates the percentage of cases among the feasible combinations of parameter choices, when N_0 , a_{ii} , and a_{ij} are fixed, for which having country $i > 0$ give up its drugs to country 0 is the Pareto-optimal Nash equilibrium. For example, when $N_0 = 10^5$, $a_{ii} = a_{00}/10$, and $a_{ij} = a_{ii}/10$, there are a total of 108 feasible combinations of N_i , K , and a_{00} , out of which 108 cases are such that having country $i > 0$ give up its drugs to country 0 is the Pareto-optimal Nash equilibrium. Therefore, the corresponding cell in Table 2 is 100% (= 108/108). It is clear from Table 2 that country i is more likely to give up its drugs in the Pareto-optimal Nash equilibrium when the

Table 2. Percentage of “giving up” equilibria.

N_0	a_{ij}	a_{ii}		
		a_{00} (%)	$a_{00}/5$ (%)	$a_{00}/10$ (%)
10^5	a_{ii}	0	81.48	99.07
	$a_{ii}/5$	0	97.22	100
	$a_{ii}/10$	0	97.22	100
	$a_{ii}/100$	0	99.07	100
10^6	a_{ii}	0	59.09	80.68
	$a_{ii}/5$	0	78.41	81.82
	$a_{ii}/10$	0	78.41	81.82
	$a_{ii}/100$	0	76.14	81.82
10^7	a_{ii}	0	25.00	55.77
	$a_{ii}/5$	0	48.08	69.23
	$a_{ii}/10$	0	51.92	57.69
	$a_{ii}/100$	0	51.92	36.54

country's a_{ii} is smaller than a_{00} , a_{ij} is smaller than a_{ii} , and N_0 is small.

We also compute the optimal solutions of the centralized planner's problem (11). In 1,366 out of the 1,390 cases where the Pareto-optimal Nash equilibrium has all drug-possessing countries keep their drugs for themselves, the centralized optimal solution corresponds to allocating all existing drugs to country 0. In the remaining 24 cases, it is optimal for the central planner to allocate drugs according to their initial endowment (i.e., each country uses the drugs it has). We further compare results from the Pareto-optimal Nash equilibrium with the centralized optimal solution. To that end, we measure the ratio of the total expected number of infectives in the world for the decentralized case over that of the centralized solution. The geometric mean of the 1,390 ratios is 1.134. That is, on average, the Pareto-optimal Nash equilibrium is 13.4% worse than the centralized optimal solution. The worst 20 out of the 1,390 cases have ratios above 2 (i.e., the decentralized allocation of drugs can lead to more than doubling the number of infectives compared to the central planner decision).

In short, our numerical study suggests that the decentralized equilibrium often leads to the optimal drug allocation of the central planner (in about half of the tested cases). When this is not the case, however, decentralized decision making seems to have a significant negative impact on the control of the pandemic.

5.2. Entire Time Horizon of the Pandemic

So far we have considered the two first periods of a multidimensional Reed-Frost model. Even though the main focus of this paper is on deriving analytical results and insights into the drug allocation problem, we provide preliminary simulation results for the entire time horizon, i.e., when considering the final sizes of the epidemic in multiple countries. Specifically, “final size” refers to the total number of infectives during the entire horizon of the epidemic. The main purpose of the simulation is to study the equilibrium

behavior when the objective of each country is the final size of the epidemics in the corresponding country. Note, however, that considering the entire time horizon of the epidemic is also problematic. Existing mathematical models and simulations of epidemics assume that transmission probabilities remain constant over the entire time horizon. However, this ignores important population behaviors such as social distancing or border closing that would certainly occur in case of a global pandemic.

We consider a world with three countries, $\{0, 1, 2\}$. Each country has a population 10^6 . We choose the transmission probability parameter $a_{00} = 2 \times 10^{-6}$. Similar to before, roughly speaking, if country 0 is isolated from the rest of the world, its basic reproductive rate is 2. The dynamics of the system is described by Equations (2)–(3). If we are given the initial allocation of the drugs, which will be explained in the next paragraph, we can simulate the number of susceptibles, infectives, and removals in the three countries over time. To speed up the simulation, we approximate the Binomial distribution with the corresponding Poisson distribution when the mean of the Binomial distribution is less than 0.02, and with Normal approximations otherwise. In each simulated trajectory, we stop until $y_i^t + \bar{y}_i^t = 0$ and obtain the final size of the epidemic z_i^∞ . The expected number of susceptibles in each time period t can be estimated by the average $x_i^t + \bar{x}_i^t$ across the many trajectories. The expected value of the total removal is estimated from the average z_i^∞ among the simulation trials. For simplicity, we call the expected value of the susceptible final size as the “final size.”

Assuming that countries $i = 1, 2$ are identical, we vary the choices of a_{ii} and K_i such that a_{ii} takes values as a_{00} , $0.9a_{00}$, $0.5a_{00}$, $0.3a_{00}$, $0.1a_{00}$; and K_i takes values 0.1, 0.3, and 0.5 of the country 0’s population size N_0 . We also set the cross country transmission probability $a_{ij} = a_{0i} = a_{i0} = a_{ii}/10$ for each case. These parameters describe symmetric network structures. Online Appendix K considers “chain” network structures in which country 2 and country 0 are not directly connected. The findings are similar, except that the two-period model has less predictive power if used as an approximation of the entire horizon model.

In total, we tested 18 combinations of model parameters (see Table 3). For each set combination of model parameters, we mutate the amount of drugs that countries 1 and 2 give up to country 0 through 0, 0.1, . . . , 1 of K_i corresponding to a total of 121 cases. We simulate 10^5 trajectories in order to have an accurate estimate of the mean final size. For each set of parameters, the simulation takes more than 10 hours on a personal computer. Our numerical results indicate that a unique Pareto-dominant equilibrium exists for every one of the entire horizon cases in Table 3, consistent with what is predicted by Theorem 2 for the two-period case. Further, in all tested cases but one, the optimal decisions at the equilibrium either give up or keep all drugs, a result that is reminiscent of the result in Theorem 1.

Table 3. Equilibria: two-period vs. entire horizon models—Symmetric network structure.

a_{ii}	K_i					
	$0.5N_0$		$0.3N_0$		$0.1N_0$	
	Two periods	Entire horizon	Two periods	Entire horizon	Two periods	Entire horizon
a_{00}	Keep	Give up	Keep	Keep	Keep	Keep
$0.9a_{00}$	Keep	Give up	Keep	Keep	Keep	Keep
$0.5a_{00}$	Keep	Give up*	Keep	Give up	Keep	Keep
$0.4a_{00}$	Give up	Give up	Keep	Keep	Keep	Keep
$0.3a_{00}$	Give up	Give up	Give up	Give up	Keep	Give up
$0.2a_{00}$	Give up	Give up	Give up	Give up	Give up	Give up

*In this equilibrium Countries $i = 1, 2$ give up $0.8K_i$ or $0.4N_0$ of drugs to country 0.

Table 3 compares the Pareto-optimal Nash equilibrium for the two-period and entire time horizon cases. In all tested cases, it turns out that countries 1 and 2 make similar allocation decisions at the equilibrium. In Table 3, as well as later in the paper, we denote “Keep” to represent the action profile in which both countries keep the drugs for their own use. Similarly, “Give up” means that both countries give up their drugs to country 0. In particular, the “Keep” and “Give up” in Table 3 indicate the corresponding action profile in the Pareto-optimal Nash equilibrium. According to Table 3, when the two-period model predicts that countries 1 and 2 give up their drugs to country 0, the same result holds at the equilibrium of the entire time horizon. On the other hand, countries sometimes keep their drugs with two periods, but give them up when considering the entire time horizon. This suggests that myopic decision makers would act more selfishly than they would if they considered the final size of the pandemic.

Finally, we compare in Table 4 the final sizes corresponding to the different drug allocations at equilibrium as reported in Table 3, for the finite and entire horizon cases. In most cases, the final sizes are equal because the equilibria from the two models are the same. When drug capacity K_i is very small ($0.1N_0$), or country i ’s transmission rate a_{ii} is very close to a_{00} , the final sizes according to different allocation policies are still similar, reflecting the fact that it does not matter very much where the drugs are used. When a_{ii} is smaller than a_{00} ($a_{ii} \leq 0.5a_{00}$) and K_i is not too small ($K_i \geq 0.3N_0$), “Give up” and “Keep” have much more profound impact on the final size of the epidemic. For example, when $a_{ii} = 0.5a_{00}$ and $K_i = 0.5N_0$, the final size of “Keep” (dominant equilibrium from the two-period model) is almost 30 times the final size of the allocation where both countries give up $0.8K_i$ to country 0 (dominant equilibrium from the entire horizon model). In this case, a myopic decision maker would act more selfishly, which, in turn, would dramatically increase the number of infectives in the world.

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at <http://journals.informs.org/>.

Table 4. Final sizes according to the dominant equilibria from the two-period model vs. the entire horizon model.

a_{ii}	K_i					
	$0.5N_0$		$0.3N_0$		$0.1N_0$	
	Two periods	Entire horizon	Two periods	Entire horizon	Two periods	Entire horizon
a_{00}	1,787,423	1,001,549	2,038,872	2,038,872	2,240,185	2,240,185
$0.9a_{00}$	1,628,703	840,109	1,909,971	1,909,971	2,124,970	2,124,970
$0.5a_{00}$	923,400	32,999	1,059,902	428,662	1,253,372	1,253,372
$0.4a_{00}$	8	8	881,695	149,608	982,405	982,405
$0.3a_{00}$	3.8	3.8	81,349	81,349	809,452	632,086
$0.2a_{00}$	3.4	3.4	64,863	64,863	545,780	545,780

6. Conclusions

Recent epidemiological studies have suggested that the prophylactic use of antiviral drugs or imperfect vaccines could slow the spread of an influenza epidemic. When existing drug stockpiles are scattered in different countries, this situation gives rise to a game where each country seeks to protect its own population. More generally, our paper provides the first insights as to how self-interested countries would actually use their antiviral drugs to cover another population if an epidemic were to appear in the world.

Existing studies of epidemic containment make extensive use of simulations. In the presence of strategic players, this approach is very hard to generalize. For our purpose, we resorted to analytical techniques and developed a tractable multivariate Reed-Frost model where countries allocate drugs according to their narrow self-interests.

The drug allocation game may evoke a prisoners’ dilemma, which would then predict that all countries keep their drugs for their own population. In fact, giving up drugs to another country could be extremely difficult to justify to the public. Nevertheless, when considering epidemic dynamics and related uncertainties, our results suggest that countries might reach an equilibrium where at least some countries give up all their drug stockpiles to contain the pandemic. We consider the very beginning of a pandemic, but preliminary simulations suggest that this insight holds for the entire time horizon, with the difference that countries give up their drugs in a broader range of situations (i.e., for more sets of parameters) than for the two-period case. Further, when countries have large populations and when the disease spreads quickly, implementing the optimal drug allocation of a central planner can reduce the number of infectives compared to the decentralized equilibrium in each country. In other words, an agreement between the different decision makers could reduce the scale of the epidemic in all countries. Surprisingly, many of our results hold for different objective functions (expected value and probability of no infection), suggesting that our findings are relatively robust to varied model assumptions.

Our model simultaneously captures three critical sources of uncertainty: the number of initial infections, the spread of the disease, and drug efficacy. Future work could analyze more refined versions of the model. For instance, a fourth

source of uncertainty comprises the number of initial infectives that receive treatment when drugs are distributed (at time $t = 0$). We assume for simplicity that none of these initially infected individuals are treated, although our results hold when this number is deterministic. An extension of our work could instead consider treating this number as a random variable. More general research directions include explicitly describing the distribution of drugs within and between countries (with lead times, for instance) that also consider logistic constraints. These extensions constitute important research opportunities, although thus far we have found these problems to be analytically demanding. Another extension is to model a longer time horizon and multiperiod decision processes. Without closed-form expressions, even a three-period model is very hard to evaluate, considering the high dimensionality of the problem (e.g., each decision maker needs to consider $X'_i, \bar{X}'_i, Y'_i,$ and \bar{Y}'_i for all countries $i \geq 0$, that is, $2(m + 1)$ dimensions, at each time epoch). In our ongoing work, the framework and findings of this paper help us derive approximations for a multiperiod dynamic game that generalizes the model. These heuristics should help us numerically test and refine the insights of this paper for richer models.

So far, we have focused on modelling the onset of an epidemic. We consider the first two periods only to mimic the epidemics exponential growth in its early stage, while keeping the model tractable. When the whole course of the pandemic is considered, and with the assumption that the transmission parameters remain constant over time, a pertinent objective for each country is to minimize the total expected number of infected people. This would, however, require the assumption that population’s behaviors do not change significantly and that countries do not close their borders during the pandemic. Nonetheless, our numerical studies suggest that our findings for two periods hold, in general, when considering the entire time horizon. On the other hand, such a game is in most cases intractable for analytical results, especially when we consider different sources of uncertainty as in our current model. However, for deterministic SIR models, each country’s response curve could be inferred from the solution of a system of equations similar to Equation (2.4.7) in Daley and Gani (2001). In our view, this constitutes the most promising research direction.

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at <http://journals.informs.org/>.

Our results also raise the intriguing question of how imperfect information affects the drug allocation game. Countries may not disclose their drug stockpiles or their actions. Country 0 may not even reveal the existence of a first infective. These extensions would help evaluate the value of sharing information in case of a pandemic. Furthermore, we assume that the model parameters (especially transmission probabilities a_{ij}) are common knowledge, which is necessary for analytical tractability. In practice, these parameters may be estimated, and may also constitute private information held by each country. It would then be interesting to determine mechanisms that give countries incentives to truthfully report their a_{ij} s. Also, for political reasons, countries might have a privately known propensity to keep drugs for themselves. This situation has been studied in the arms race and negotiation literature (see Baliga and Sjostrom 2004) and constitutes a valuable research direction for the drug allocation game we study in this paper.

Finally, our paper is the first to study the decentralized allocation of resources to control a pandemic. This framework appears very relevant for the containment of epidemics in other contexts. For example in a recent working paper, Chick et al. (2009) study coordination contracts for a game between governments that purchase vaccines in order to achieve an efficient cost-benefit trade-off in their respective countries.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org/>.

Endnotes

1. The H5N1 influenza A virus is the virus that is the subject of greatest concern for a pandemic at present.
2. These simulations allow great precision in modeling the spatiotemporal evolution of the epidemic as well as different drug or vaccine distribution policies (e.g., ring policies or vaccination of children). The resulting estimates of the epidemic size and the probability of extinction are therefore fairly accurate. On the other hand, such detailed dynamics are analytically intractable, making structural results hard to obtain. In addition, simulations are inadequate to analyze equilibria resulting from drug distributions made by strategic players, which is the main focus of this paper. Thus, this paper focuses on deriving insights and structural results based on simplified models.
3. Given this particular structure of the game, our problem becomes similar to the seminal model proposed by Bergstrom et al. (1986) on the private provision of public goods. Our model, however, violates important assumptions made in Bergstrom et al. (1986). As a result, the analysis of Bergstrom et al. (1986) does not apply to our analysis.
4. Generally speaking, the set of Nash equilibria is a subset of correlated equilibria, and the “best” correlated equi-

librium may improve upon Nash. Unfortunately, the set of Nash and correlated equilibria have the same bounds for supermodular games, as indicated by Milgrom and Roberts (1990). This means that correlated equilibria do not improve in our problem.

Acknowledgments

The authors appreciate useful discussions with and feedback from Bill Gimson, Kimberly Lane, and Fangjun Zhou at the Centers for Disease Control and Prevention, Edward Kaplan at Yale University, Duncan Simester at Massachusetts Institute of Technology, Padmaja Patnaik at the Family Health International (Epidemiology), Karin Yeatts at the Department of Epidemiology at the University of North Carolina, and seminar participants at Duke University.

References

- Anonymous. 2005. Preparing for an influenza pandemic. *The Economist* (Sep. 24) 95–96.
- Baliga, S., T. Sjostrom. 2004. Arms races and negotiations. *Rev. Econom. Stud.* **71**(2) 351–369.
- Ball, F. G. 1991. Dynamic population epidemic models. *Math. Biosciences* **107** 299–324.
- Bauch, C. T., D. J. D. Earn. 2004. Vaccination and the theory of games. *Proc. National Acad. Sci.* **101**(36) 13391–13394.
- Bergstrom, T., L. Blume, H. Varian. 1986. On the private provision of public goods. *J. Public Econom.* **29** 25–49.
- Brandeau, M. L. 2004. *Allocating Resources to Control Infectious Diseases: Operations Research in Health Care*, Chapter 17. Kluwer Academic Publishers, 443–465.
- Brandeau, M. L., G. S. Zaric, A. Richter. 2003. Resource allocation for control of infectious diseases in multiple independent populations: Beyond cost-effectiveness analysis. *J. Health Econom.* **22**(4) 575–598.
- Carlsson, H., E. Van Damme. 1993. Global games and equilibrium selection. *Econometrica* **61** 989–1018.
- Chick, S. E., H. Mamani, D. Simchi-Levi. 2008. Supply chain coordination and influenza vaccination. *Oper. Res.* **56**(6) 1493–1506.
- Chick, S. E., H. Mamani, D. Simchi-Levi. 2009. Influenza vaccine supply chain with multiple agencies. Working paper, University of Washington, Seattle.
- Cho, S.-H. 2009. The optimal composition of influenza vaccines subject to random production yields. *Manufacturing Service Oper. Management*. ePub ahead of print September 25, <http://msom.journal.informs.org/cgi/content/abstract/msom.1090.0271v1>.
- Daley, D. J., J. Gani. 2001. *Epidemic Modelling: An Introduction*. Cambridge University Press, Cambridge, UK.
- Deo, S., C. J. Corbett. 2009. Cournot competition under yield uncertainty: The case of the U.S. influenza vaccine market. *Manufacturing Service Oper. Management* **11**(4) 563–576.
- Diekmann, O., J. Heesterbeek. 2000. *Mathematical Epidemiology of Infectious Diseases*. Wiley, New York.
- Ferguson, N. M., S. Mallett, H. Jackson, N. Roberts, P. Ward. 2003. A population-dynamic model for evaluating the potential spread of drug-resistant influenza virus infections during community-based use of antivirals. *J. Antimicrobial Chemotherapy* **51** 977–990.
- Ferguson, N. M., D. Cummings, S. Cauchemez, C. Fraser, S. Riley, A. Meeyai, S. Iamsrithaworn, D. S. Burke. 2005. Strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature* **437** 209–214.
- Frankel, D. M., S. Morris, A. Pauzner. 2003. Equilibrium selection in global games with strategic complementarities. *J. Econom. Theory* **108** 1–44.

- Germann, T. C., K. Kadau, I. M. Longini, C. A. Macken. 2006. Mitigation strategies for pandemic influenza in the United States. *Proc. National Acad. Sci.* **103**(15) 5935–5940.
- Harsanyi, J. C., R. Selten. 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press, Cambridge, MA.
- Hill, A. N., I. M. Longini. 2003. The critical vaccination fraction for heterogeneous epidemic models. *Math. Biosciences* **181** 85–106.
- Kornish, L. J., R. L. Keeney. 2008. Repeated commit-or-defer decisions with a deadline: The influenza vaccine composition. *Oper. Res.* **56**(3) 527–541.
- Lau, J., H. Tsui, J. H. Kim, S. Griffiths. 2006. Perceptions about status and modes of H5N1 transmission and associations with immediate behavioral responses in the Hong Kong general population. *Preventive Medicine* **43**(5) 406–410.
- Longini, I. M., A. Nizam, S. Xu, K. Ungchusak, W. Hanshaworakul, D. A. Cummings, M. E. Halloran. 2005. Containing pandemic influenza at the source. *Science* **309**(12) 1083–1087.
- Marx, L. M., S. A. Matthews. 2000. Dynamic voluntary contribution to a public project. *Rev. Econom. Stud.* **67** 327–358.
- Milgrom, P., J. Roberts. 1990. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* **58** 1255–1277.
- Mills, C. E., J. M. Robins, C. T. Bergstrom, M. Lipsitch. 2006. Pandemic influenza: Risk of multiple introductions and the need to prepare for them. *PLoS Medicine* 0769–0773.
- Monto, A. S. 2005. The threat of an avian influenza pandemic. *Nature* **352**(4) 323–325.
- Myerson, R. B. 1991. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA.
- Philipson, T. 2004. Economic epidemiology and infectious disease. J. Newhouse, T. Culyer, eds. *Handbook of Health Economics*, Chapter 17. North-Holland, New York, 443–465.
- Rosenthal, E. 2005. Better planning is needed for flu drugs, experts say. *The New York Times* (October 19) 8.
- VanHuyck, J., R. C. Battalio, R. O. Beil. 1990. Tacit coordination games, strategic uncertainty, and coordination failure. *Amer. Econom. Rev.* **80**(1) 234–248.
- VanHuyck, J., R. C. Battalio, R. O. Beil. 1991. Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *Quart. J. Econom.* **106**(3) 885–910.
- Ward, P., I. Small, J. Smith, P. Suter, R. Dutkowski. 2005. Oseltamivir (Tamiflu (R)) and its potential for use in the event of an influenza pandemic. *J. Antimicrobial Chemotherapy* **55** i5–i21.
- Wu, J. T., L. M. Wein, A. S. Perelson. 2005. Optimization of influenza vaccine selection. *Oper. Res.* **53**(3) 456–476.
- Yang, Y., I. M. Longini Jr., M. E. Halloran. 2006. Design and evaluation of prophylactic interventions using infectious disease incidence data from close contact groups. *J. Roy. Statist. Soc. Ser. C (Appl. Statist.)* **55** 317–330.