DISCUSSION PAPER

September 2011 ■ RFF DP 11-41

Policy Response to Pandemic Influenza

The Value of Collective Action

Georgiy Bobashev, Maureen Cropper, Joshua Epstein, Michael Goedecke, Stephen Hutton, and Mead Over

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

RESOURCES FOR THE FUTURE

Policy Response to Pandemic Influenza: The Value of Collective Action

Georgiy Bobashev, Maureen Cropper, Joshua Epstein, Michael Goedecke, Stephen Hutton, and Mead Over

Abstract

This paper examines positive externalities and complementarities between countries in the use of antiviral pharmaceuticals to mitigate pandemic influenza. It demonstrates the presence of treatment externalities in simple SIR (susceptible-infectious-recovered) models and simulations of a Global Epidemiological Model. In these simulations, the pandemic spreads from city to city through the international airline network and from cities to rural areas through ground transport. While most treatment benefits are private, spillovers suggest that it is in the self-interest of high-income countries to pay for some antiviral treatment in low-income countries. The most cost-effective policy is to donate doses to the country where the outbreak originates; however, donating doses to low-income countries in proportion to their populations may also be cost-effective. These results depend on the transmissibility of the flu strain, its start date, the efficacy of antivirals in reducing transmissibility, and the proportion of infectious people who can be identified and treated.

Key Words: pandemic influenza, disease control externalities

JEL Classification Numbers: I18, C63, D62

Contents

Policy Response to Pandemic Influenza: The Value of Collective Action

Georgiy Bobashev, Maureen Cropper, Joshua Epstein, Michael Goedecke, Stephen Hutton, and Mead Over*

1. Introduction

In the 20th century, the world experienced three influenza pandemics that had significant economic costs and caused millions of deaths and illnesses.[1] The deadliest of these, the 1918 pandemic, killed an estimated 3 percent of the world's population despite being slow to spread in an age before commercial air travel. With today's global transportation network, an outbreak of influenza could quickly reach pandemic proportions. The 2009 H1N1 "swine flu" pandemic and the SARS outbreak of 2003 remind us of this risk.

The world's response to these outbreaks has been characterized by relatively little cooperation and coordination between countries, with wealthy nations stockpiling doses of vaccines or antivirals for their own citizens before considering treatment in other countries. Whether such an inwardly focused policy is optimal depends on the nature and magnitude of externalities in treating pandemic flu. How much do policies that slow the spread of the flu in one country reduce attack rates in other countries? Does treatment of infected persons in one country increase the marginal benefits of treatment policies in other countries? The overarching goal of this paper is to answer both questions in a realistic model of the spread of influenza through the global air transport network.

By simulating the impact of control strategies in a global epidemiological model, we are able to address two questions regarding international cooperation to mitigate a pandemic:

---

* Georgiy Bobashev and Michael Goedecke, RTI International; Maureen Cropper, University of Maryland and Resources for the Future; Joshua Epstein, Johns Hopkins University; Stephen Hutton, World Bank; Mead Over, Center for Global Development. This work was supported in part by the Pilot Studies of Modeling of Infectious Disease Agents Study (MIDAS) cooperative agreement from NIGMS (1 U01 GM0700698). We thank the World Bank Research Board and the KCP Trust Fund for research support and seminar participants at the Center for Global Development and World Bank for useful comments.

[1] These were the 1918,1957, and 1968 pandemics.

- Is it cost-effective for wealthy nations to fund the purchase and distribution of antivirals in low-income countries to slow the spread of the pandemic?
- What global allocation rules are most effective in reducing attack rates?

When treatment with antiviral drugs alone can contain a potential pandemic originating in a low-income country, wealthy nations have an obvious incentive to pay for them. But what about the case where a pandemic cannot be detected early enough or treatment is not effective enough for containment to be possible? Pandemic epidemiology involves two types of externalities that suggest that it might be in the self-interest of wealthy countries to fund treatment in such a scheme. First, the treatment policy in one country will affect the rate at which the pandemic spreads to other countries, so treatment provides a positive externality. At the same time, the increasing marginal effectiveness of treating more people can lead to complementarities across countries. The question is, how large are these effects?

Our Approach

To investigate these issues, we use the Global Epidemiological Model (GEM) to simulate influenza pandemics under a range of conditions and antivirus treatment policies. The model divides the world into 106 Global Trade Analysis Project (GTAP) regions and models air travel among 288 cities in these regions. The flu spreads from one city to another via air travel and from cities to rural areas in each region via land travel. Within each city or rural area, the flu spreads via the SIR model, in which people transition from susceptible (S) to infectious (I) to recovered or dead (R). The model distinguishes among age groups and uses age-specific contact rates to model the spread of the flu. The probability of infection given contact in a particular region varies with latitude and season — it decreases as one moves away from the equator and is higher during the winter than during the summer.

We focus on mitigation of the spread of the flu by administering antivirals to symptomatic infectious individuals. In the absence of international cooperation, we assume that antiviral stockpiles, as a percent of population, vary with per capita income. Low-income countries are assumed to have no stockpiles. We ask whether it pays high-income countries to pay for stockpiles to be distributed in low-income countries. We compare two rules for high-income countries distributing stockpiles in low-income countries: one under which each country receives a fixed number of doses in proportion to population and another under which antivirals are allocated to the country in which the flu begins, which we assume to be a low-income country. We judge the success of a control strategy in terms of its impact on the attack rate, or

percent of the population infected, at the end of a year. Does it collectively pay high-income countries to make a donation in terms of the impact it has on their own attack rate?

With or without control policies, the percent of the population infected in each region depends on when the flu begins and the infectiousness of the flu. For example, influenza that peaks in Northern Hemisphere winter infects more people. The answers to our questions are also contingent on these parameters.

Preview of Results

The benefits from collective action in the form of influenza treatment depend on the size of treatment externalities. Treatment externalities are large if treatment can contain an outbreak in the source country (Ferguson et al. 2005; Longini et al. 2005) but smaller if the outbreak develops into a pandemic. In this case, treating infectious people in one's own country reduces the domestic attack rate but has a proportionately much smaller impact on other countries. The question is whether the cost of purchasing and distributing antivirals to other countries pays for itself in terms of reducing a country's own attack rate.

The benefits to high-income countries from paying for antivirals in low-income countries do cover the costs under "midrange" assumptions of influenza transmissibility: donation of antivirus to the country where the outbreak originates reduces the number of influenza cases in high-income countries after one year by 4.76 million cases at the cost of roughly 1.92 doses per case avoided. This donation policy is welfare enhancing for wealthy countries even when the fatality rate is zero; at any positive fatality rate the policy is even more valuable.

The paper is organized as follows. Section 2 briefly reviews the literature on policies to control pandemic flu. In section 3 we discuss the dynamics of pandemic influenza in a simple one-city SIR model in which the decision maker determines the proportion of infectious people to treat with antivirals before the pandemic begins (and hence the size of a stockpile to hold).Section 4 presents a two-city case, in which the flu spreads through travel from one city to another. We use this model to illustrate the nature of externalities and complementarities in influenza treatment. We also contrast the Nash equilibrium in treatment strategies with the cooperative treatment strategy. Section 5 presents the GEM, which we use for our simulations; section 6 presents the results of simulating pandemic flu under various antiviral stockpile assumptions. Section 7 concludes.

2. Relationship to the Literature

Models of human transmission of influenza include those with significant detail at the country level and global models of influenza transmission. Spatially detailed models of avian flu transmission at the country level have been used to compare the effectiveness of ring quarantines, vaccination strategies, and various policies for distributing antiviral drugs—for example, distributing drugs to all persons in the same school or workplace as an infected person, or, alternately, ring prophylaxis. Notable work includes models of influenza transmission in Thailand and Southeast Asia (Ferguson et al. 2005; Longini et al. 2005) and in the United States and the United Kingdom (Ferguson et al. 2006; Germann et al. 2006).

Global models, which capture the transmission of influenza through the global air transport network, have been used to examine the impact of restricting travel between countries—as a result of either government policy or private averting behavior — and traveler quarantines. These models follow the seminal work of Rvachev and Longini (1985), who used a deterministic difference equation model to simulate travel within a network of 52 cities. The model, which was used to replicate the spread of the 1968 influenza pandemic, used actual air travel data to model movement between cities. Behavior within cities was modeled more simply, assuming a uniform mixing of individuals within each city, with disease dynamics governed by a susceptible-exposed-infectious-recovered (SEIR) model. The subsequent literature (Grais et al. 2003; Hufnagel et al. 2004; Cooper et al. 2006; Colizza et al. 2007; Epstein et al. 2007) has introduced randomness in travel rates and disease transmission. It also has used a more detailed characterization of the air transport network; for example, Colizza et al. (2007) model 3,500 airports. These global models have examined the impact of air travel restrictions and the effects of local policies to reduce disease transmission. But they have not examined the economics of pandemic flu control or characterized the nature of externalities in treatment of pandemic flu.

3. Antiviral Usage in a Simple SIR Model

The simplest model of influenza spread, which underlies all of our analyses, is the SIR model. Individuals are susceptible, infectious, or recovered (immune or dead). Susceptible and infectious persons are assumed to mix uniformly. Letting $S(t)$, $I(t)$, and $R(t)$ denote the fractions of the population in each state, the changes over time in the health of the population are given by the following:

(1) $$dS(t)/dt = -\beta S(t)I(t),$$

(2) $dI(t)/dt = \beta S(t)I(t) - \delta I(t)$,

(3) $dR(t)/dt = \delta I(t)$.

β represents the probability of transmission conditional on exposure to an infected person, or the average rate of infection per susceptible person. With uniform mixing, $\beta S(t)$ represents the rate at which an infectious individual infects susceptible persons. δ is the rate at which an infected individual recovers; thus $1/\delta$ is the average duration of the disease. Disease prevalence increases over time ($dI(t)/dt > 0$) if and only if $[\beta/\delta]S(t) > 1$. Simply put, prevalence increases only if the number of infections caused by an infectious person during the time he is infectious exceeds 1 (i.e., only if he can replace himself). Thus, a necessary condition for an epidemic to begin is for $\beta/\delta > 1$. β/δ is also the average number of persons infected by an infectious person in an otherwise totally susceptible population (i.e., when $S(0) \approx 1$) and is termed the basic reproductive rate, R_0.

Figure 1 illustrates the progression of the flu for the case of $\beta = 0.3$, $\delta = 0.2$, implying that the flu lasts five days (on average) and that $R_0 = 1.5$.(Initially, $I(0) = 0.001$.) The flu spreads slowly when the stock of infectious is low, spreads more rapidly as the proportion of infectious rises, peaks when the proportion of susceptible people falls below $1/R_0$, and steadily declines thereafter. The epidemic is effectively over in less than four months, with only 0.01N new infectious cases occurring after day 104, where N is the population size. The attack rate, $R(\infty)$, is approximately 58 percent.[2]

Suppose now that some proportion of infectious people can be treated with an antiviral drug that reduces their infectiousness. Assume that this proportion must be chosen before the epidemic begins and must remain constant throughout. Let p be the proportion of new infectious people receiving antiviral treatment, and assume that this reduces the untreated rate of infection, $\beta°$:

(4) $\beta = \beta°(1-pe)$,

Where e is the proportionate reduction in $\beta°$ achieved by the antiviral (e.g., e= 0.4 implies a 40 percent reduction in the infectiousness of a sick person). Higher values of p reduce infectiousness and, hence, R_0.

[2] Technically, the attack rate is $R(\infty) - R(0)$, but we assume $R(0) \approx 0$ and so refer to $R(\infty)$ throughout.

Figure 1. Progression of Influenza in a One-City Susceptible-Infectious-Recovered (SIR) Model

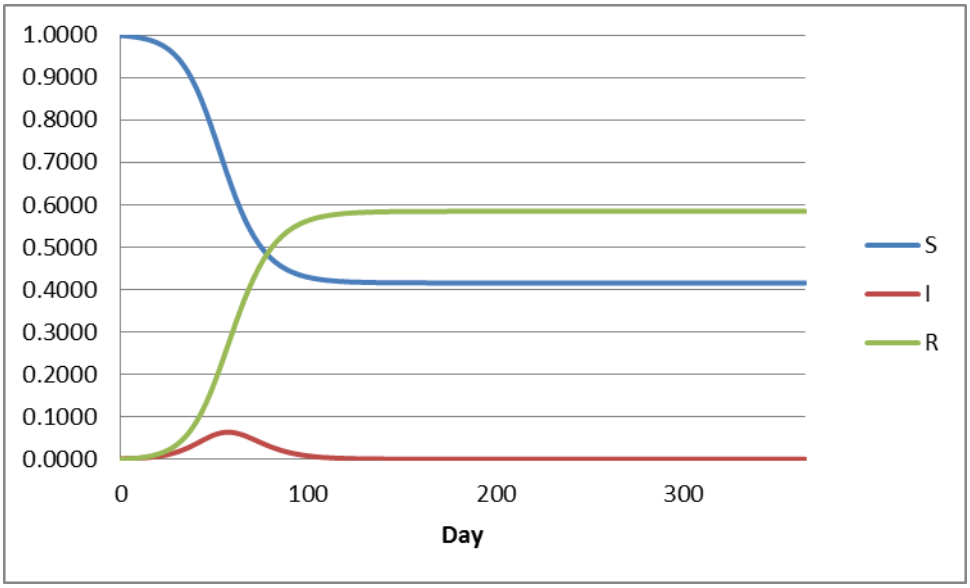

Assume that the government chooses the proportion of infectious people to treat to maximize the difference between the value of cases avoided and the costs of treatment. The benefits of treatment equal the value per case avoided, V, times the number of people who avoid infection due to treatment, $(S(\infty,p) - S(\infty,0))N$. For simplicity, we remove the $-S(\infty,0)$ term (which is a constant for any R_0 and so drops out in all derivatives) and suppress the argument of $S(\infty,p)$. Assuming that the cost per antiviral dose is constant and equal to c per case treated, the government chooses p before the epidemic begins to maximize:[3]

(5) $$F(p) = V[S(\infty)]N - cp[1-S(\infty)]N.$$

The relationship between p and the proportion of people never infected is given by the transcendental equation (Kermack and McKendrick 1927):

(6) $$ln(S(\infty)) = R_0(S(\infty)-1),$$

Where $R_0 = \beta°(1-pe)/\delta$.

[3] This assumes that the decision maker focuses on the attack rate. Antivirals also could reduce the duration of illness for those who become infected, but we ignore this benefit.

For values of R_0 that characterize pandemic flu, the p that maximizes equation (5) is a corner solution: either p=0 or p=1. First-order conditions for the maximization of (5) require p to be chosen at the point where the marginal benefits of increasing $S(\infty)$and the number of doses needed to treat infectious people equal or exceed the cost of an additional dose:

(7) $$dF(p)/dp = (V + cp)[dS(\infty)/dp] - c[1-S(\infty)] \geq 0.$$

In contrast, the second-order conditions for an interior maximum, $d^2F(p)/dp^2 < 0$, are not satisfied for pandemic flu.[4] Equation (6) implies that the proportion of people never infected, $S(\infty)$, increases as R_0 falls (i.e., as p is increased); however, for values of R_0>1.05, increasing p increases $S(\infty)$ at an increasing rate.[5] Influenza pandemics have generally been characterized by reproductive rates R_0>1.5.[6] This implies that p=0 maximizes (5) if V is low relative to c, but p=1 maximizes (5) if V is high relative to c. To illustrate, when $\beta° = 0.3$, $\delta = 0.2$, and $e = 0.2$, setting $p = 1$ is optimal if $V/c > 1.2$.

When $p = 1$, the number of doses of antivirals used equals the attack rate times the size of the population, $R(\infty)N$. We term this the antiviral stockpile used. In the example where $\beta° = 0.3$, $\delta = 0.2$, and $e = 0.2$, treatment with antivirals reduces $R(\infty)$ from 0.58 to 0.33, hence the antiviral stockpile (the total number of doses used) = $0.33N$.

In practice, the number of doses of antivirals used will be even smaller due to constraints on p. The difficulty in identifying infectious people and treating them with antivirals implies that in practice, p will be constrained to be less than 1 by the health care delivery system; formally, $p \leq p^*$. In the United States, it usually assumed that $p \leq 0.6$ (Germann et al. 2006), and the maximum feasible value of p is likely to be even lower in developing countries. If p=1 is optimal in the unconstrained problem, then $p = p^*$ and the number of antiviral doses that is used will = $p^*R(\infty)N$. In our example with $\beta° = 0.3$, $\delta = 0.2$, and $e = 0.2$, R_0 is reduced to 1.32 if p= 0.6 (rather than 1.2 when p=1). In this case, the attack rate will be $R(\infty) = 0.45$ and a stockpile for 27

---

[4] $d^2F(p)/dp^2 = (V + cp)[d^2S(\infty)/dp^2] + c[dS(\infty)/dp]$, which need not be negative.

[5] A sufficient condition for $d^2F(p)/dp^2 > 0$ is $d^2S(\infty)/dR_0^2 > 0$. Formally, $dS(\infty)/dR_0 = [S(\infty)-1]/[S(\infty)^{-1} - R_0]$and $d^2S(\infty)/dR_0^2 = [S(\infty)^{-1} - R_0]^{-1}[dS(\infty)/dR_0][2 + S(\infty)^{-2}[dS(\infty)/dR_0]]$.
This second derivative $d^2S(\infty)/dR_0^2 > 0$ provided $2 + S(\infty)^{-2}[dS(\infty)/dR_0] < 0$. Equation (6) implies that this condition is satisfied if R_0> 1.05.

[6] The R_0 for the 1957 Asian flu has been estimated at 1.8 percent (Vynnycky and Edmunds 2008). Vynnycky et al. (2007) estimate an R_0 of 2.4-4.3 for the 1918 Spanish flu.

percent of the population will be used. If $e = 0.4$ and $p = 0.6$, the attack rate will be $R(\infty) = 0.24$, and the stockpile approximating 14.4 percent of the population will be used.

4. Antiviral Usage in a Two-City Model

To illustrate the nature of externalities involved in antiviral distribution we turn to a two-city model. The two cities are linked by travel, and the flu spreads in each city according to the SIR model (see Figure 2). In this example the cities are symmetric — i.e., population (N), transmissibility of the flu, and its duration are identical in both cities, as is the rate of travel between cities. The dynamics of the spread of the flu are given by equations (8) – (11) for city A:

(8) $dS_A/dt = -\beta_A S_A I_A - S_{MA} + S_{MB}$,

(9) $dI_A/dt = \beta_A S_A I_A - I_{MA} + I_{MB} - \delta I_A$,

(10) $dR_A/dt = \delta I_A$,

(11) $\beta_A = \beta^0_A(1-p_A e)$,

Where S_A, I_A, and R_A are shorthand for $S_A(t)$, $I_A(t)$, and $R_A(t)$, respectively, and "M" subscripts denote migrants. Equations for city B are defined analogously. We assume that a fraction α of infectious people travel between cities and, to keep the population constant, that the same number of susceptible people also travel.

Figure 2. Spread of the Flu in a Two-City Susceptible-Infectious-Recovered Model

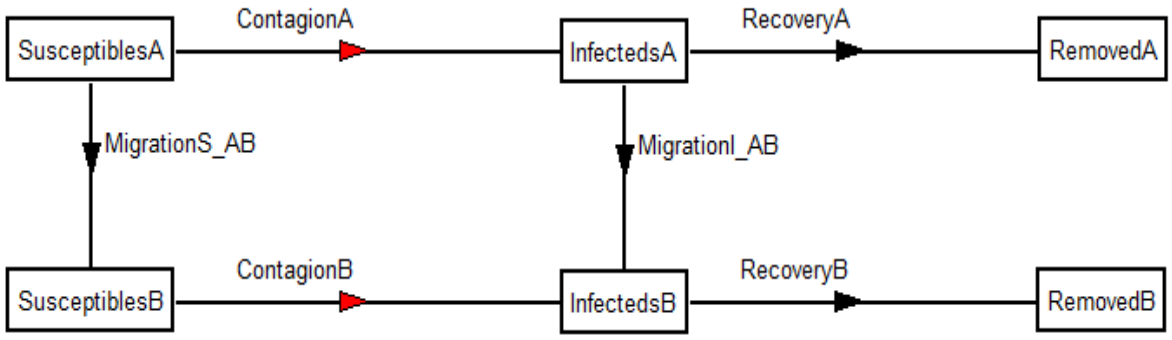

Figure 3 illustrates disease dynamics for the two cities assuming that the flu starts in city A and moves to city B ($I_A(0) = 0.001$, $I_B(0) = 0$). The figure assumes that neither city uses

antivirals (p_A=p_B=0) and that $\beta^0 = 0.3$,$\delta = 0.2$, and α=0.01. The number of infectious people peaks first in the outbreak city (city A), but the attack rates are approximately equal (58 percent) in the two cities.

Figure 3. Progression of Influenza in a Two-City Susceptible-Infectious-Recovered Model

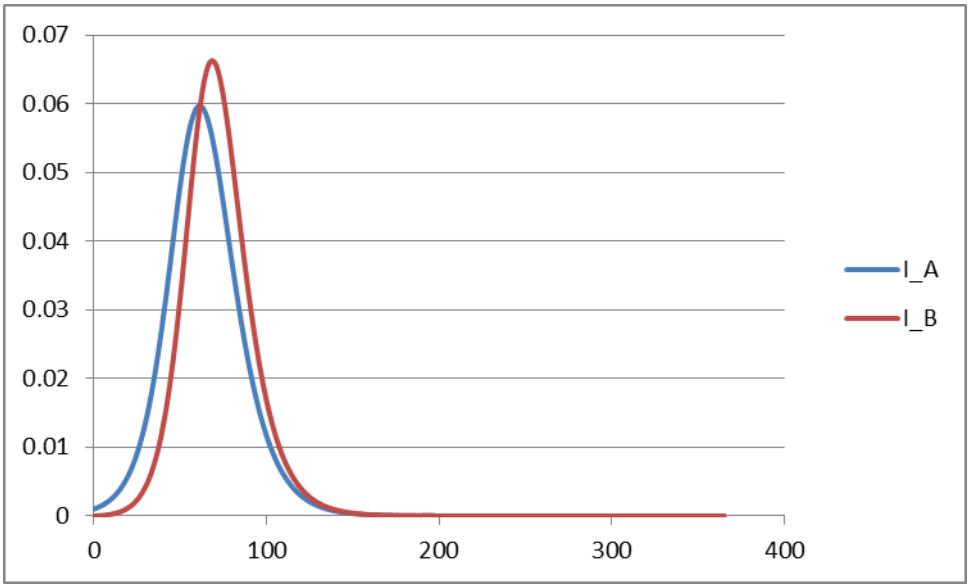

How large are the externalities associated with the use of antivirals? If city A uses antivirals, how much does it lower the attack rate in city B, $R_B(\infty)$, and vice versa? Table 1 illustrates the nature of the externalities from antiviral treatment. Specifically, the table shows the attack rate in each city as a function of the proportion of infectious treated by each of the two cities. We maintain same parameter values as in the single city case ($p^* = 0.6$, $e = 0.4$, $\beta^0 = 0.3$, and $\delta = 0.2$), hence each city chooses to treat either 60 percent of its infectious or none.

Three points about Table 1 deserve emphasis. First, when only one city treats its infectious, the effectiveness of its treatment is less than in the one-city case where there is no travel. When city A treats 60 percent of its infectious in isolation, the attack rate is reduced from 58 percent to 24 percent. In the two-city case, the attack rate is reduced to 30 percent. People from city B continue to re-infect city A, making city A's treatment of its infectious less effective.

Second, when only one city treats its infectious, the external benefits to the no-treatment city are small: when city B has no stockpile, city A's treatment of its infectious reduces city B's attack rate by only 1.6 percentage points.

Third, when both cities treat, the effectiveness of each city's stockpile is greater than when only one city employs a treatment strategy. In this case, the attack rate is reduced from about 58 percent to 24 percent in each city, as in the single city case. In effect, the two cities have become a single city, employing the same treatment as in the single-city case.

The magnitude of the externalities and complementarities observed in Table 1 are affected by the transmissibility of the flu (β^0/δ), the proportion of infectious who can be treated (p^*), and the effectiveness of treatment (e). The benefits to city B of antivirals in city A are greater the less transmissible is the flu, as are the complementarities of treatment (see Table 1). Treatment complementarities are also greater the higher are the proportion of infectious who can be treated (p^*) and the greater is the efficacy of treatment (e) (see Appendix A).

Table 1. Externalities in Antiviral Treatment, Two-City Model

β°	p_A	p_B	Attack rate in city A	Attack rate in city B
0.3	0	0	0.588	0.579
	0.6	0	0.301	0.563
	0	0.6	0.565	0.296
	0.6	0.6	0.240	0.235
0.35	0	0	0.719	0.707
	0.6	0	0.475	0.705
	0	0.6	0.709	0.468
	0.6	0.6	0.456	0.449
0.25	0	0	0.376	0.370
	0.6	0	0.111	0.326
	0	0.6	0.329	0.107
	0.6	0.6	0.012	0.005

Note: β^0= untreated rate of infection; p_A= proportion of new infectious people receiving antiviral treatment in city A; p_B= proportion of new infectious people receiving antiviral treatment in city B.

Nash Equilibrium

An important question is what treatment strategy each city will choose, given the behavior of the other city. To examine the Nash equilibrium in treatment strategies, we assume that each city selects the proportion of infectious it will treat so as to maximize the objective function in (5), modified such that $S_i(\infty)$, $i = A, B$, depends on the proportion treated in both cities. We maintain our previous assumptions that $\beta^\circ = 0.3$, $\delta = 0.2$, $e = 0.4$, $\alpha = 0.01$, and $p^* = 0.6$. We begin with a case in which the value of avoiding a case of the flu is 10 times the cost of

a course of treatment (specifically, V=100 and c=10) in both cities. This is a ~~"h~~high-value" scenario that would result in each city choosing $p = p^*$ if the city operated in isolation.

Figure 4 displays each city's best response function in the high-value scenario. The Nash equilibrium is p_A=p_B=0.6, which is also the treatment strategy that maximizes the sum of cases avoided minus treatment costs for the two cities combined, and is thus the efficient outcome.

Figure 4. Nash Equilibrium, High-Value Scenario

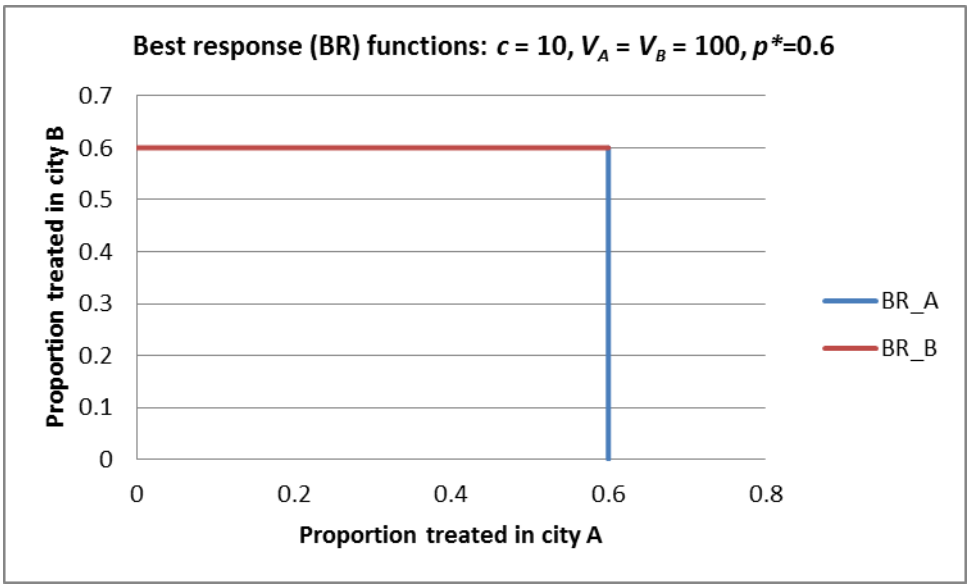

Suppose now that city A is a low-income city, where the value of avoiding a case is less than the cost of treatment (V_A = 5, c_A = 10). The value of avoiding a case remains 10 times the cost in city B (V_B = 100, c_B = 10). The best response functions for the low-income city and high-income city are pictured in Figure 5. The best response for city A depends on the value of p_B; at low values, city A will choose p_A = 0, but at high values, city A will choose p_A= 0.6, and the Nash equilibrium remains at p_A=p_B=0.6, which is still socially efficient

Figure 5. Nash Equilibrium, Lower-Value Scenario in City A

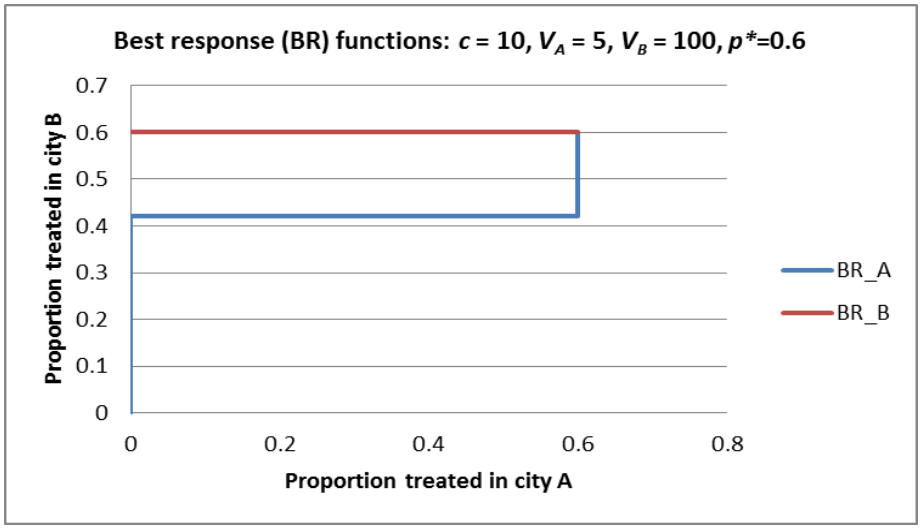

Now suppose that the value in city A is even lower ($V_A = 3$), while the cost of treatment remains constant. The best responses for each city are pictured in Figure 6. The Nash equilibrium is $p_A= 0, p_B=0.6$, whereas the treatment strategy that maximizes the sum of cases avoided minus treatment costs for the two cities combined is still $p_A= p_B= 0.6$, so the Nash equilibrium is inefficient. It is also the case that city B benefits by paying for antivirals sufficient to treat 60 percent of infectious people in city A. The cost of the treatment ($10*(0.6)*0.235N$) is less than the value of the resulting reduction in city B's influenza cases ($100*0.061N$), so B's purchase of doses for A is a pareto improvement.

Figure 6. Nash Equilibrium, Lowest-Value Scenario in City A

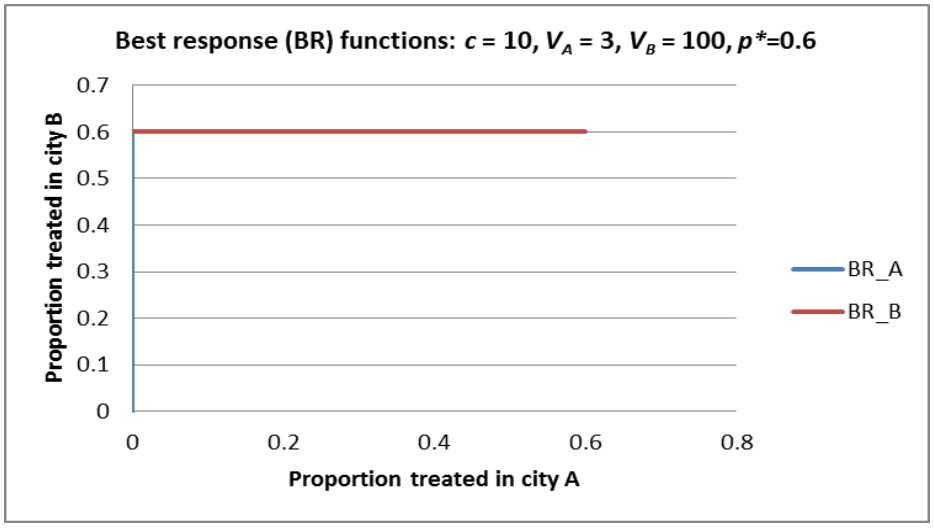

Whether city B benefits by paying for antiviral treatment in city A depends on the transmissibility of the flu, the efficacy of treatment, and the maximum proportion of infectious people who can be treated. A policy of funding antiviral treatment in City A does not pass the benefit–cost test for a more readily transmissible flu ($\beta^\circ = 0.35$) when $V_B = 100$ and $c = 10$, and the benefits are lower than in the Figure 6 scenario when e=0.3 and p^*=0.5 (Appendix A), even though they remain positive. Given the sensitivity of the results to particular parameter values, we move to a model with a higher degree of descriptive realism.

5. The Global Epidemiological Model

While simple one- and two-city models can demonstrate the core qualitative properties of categorical epidemiological models, it is difficult to draw policy conclusions from these models. To simulate realistic treatment policies, we have expanded and modified the Global Epidemiological Model developed by Bobashev and Goedecke (Hajdin et al. 2009). The GEM is a discrete time, stochastic SEIR model designed to simulate the spread of a global pandemic. The model is based on the work of Rvachev and Longini (1985) and on the epidemic model of Baroyan et al. (1981). The GEM extends these models by adding stochastic components, disease interventions, and a more detailed population structure.[7]

Network Structure

The GEM tracks the spread of a pandemic through a network of 288 cities and 101 rural areas. The world is divided into 106 GTAP regions, each of which is an individual country or a group of countries (see Appendix C).[8] Every region contains at least one city, and the network is fully connected.[9] Cities include the 130 most populous in the world and the 130 with the highest average airport flows per day.[10] A rural area—any area that the model does not explicitly define as urban—exists for every region except small islands and city states, and collectively account

[7] Though the GEM is stochastic, the simulations reported here are run using a deterministic version of the model. The attack rate (our primary outcome variable) varied little across runs of the stochastic version of the model.

[8] Roughly 86 percent of the world's population lives in a country that is its own region. The rest live in aggregate regions like "Rest of Caribbean," "Rest of Eastern Africa," and "Rest of Western Asia."

[9] There is a path from each node in the network to every other node.

[10] As a result, we end up with many cities in high-income countries, and relatively fewer cities in low-income countries. Runs of the model excluding some high-income country cities produced results similar to those reported in this paper.

for 84 percent of the total 6.4 billion world population. Rural areas are treated the same as cities in the model except that they are connected only to the cities within the same region.

The cities are connected to rural areas by a daily flow of people using ground transportation, and cities are connected to each other by airline flows. The model assumes that 1 percent of each urban population travels to its associated rural area each day, and an equal number travel from the rural area to the city.[11] Although this assumption is somewhat arbitrary, model results are not very sensitive to changes in travel rates.[12] Cities are connected through 7,668 city-pair connections with a total global daily air movement of 5,035,664 people per day. The model randomly selects travelers from the available pool of susceptible, exposed, asymptomatic infectious, and recovered cases; symptomatic infectious cases do not travel. This means that the pool of travelers on average is exposed or asymptomatic infectious in the same proportions as the general population.

Travel and population data come from a variety of sources. Data on airline flows are drawn from Official Airline Guide statistics on flight schedules provided by Guimerà et al. (2005). Numbers of passengers are based on the daily seat capacity between each pair of cities. The raw data track only direct flights, but we assume a proportion of travelers take two-leg, indirect flights.[13]

Data on regional populations come from the World Bank development data portal. We draw data on urban populations primarily from the United Nations' (2007) World Urbanization Prospects estimates, which use broad definitions of urban areas.[14] We calculate rural populations by deducting urban populations from regional populations.

---

[11] Attempts to estimate rates of real-world land travel using traffic count data encountered several problems. First, while data were available in many high-income countries, data in developing countries were sparse. Second, even the high-quality developing-country data could be extremely misleading, as data on traffic count often include significant numbers of commuters from areas just outside urban boundaries, leading to estimates that could be as high as 15 percent of the value of the urban population. Third, it was not possible to find reliable data for non-car-based travel.

[12] Running the model with doubled and halved travel rates had no significant impact on attack rates by the end of the pandemic.

[13] We use a dataset of 10 percent of domestic U.S. airline ticket coupons for 2004 to estimate the proportion of travelers on any city pair who travel through a connecting airport, rather than directly. These estimates are then applied to the cities in the GEM. On average, we assume 68.2 percent of tickets are direct and 31.8 percent have two legs (though this varies across each city pair in the model).

[14] Some cities are urban agglomerations of multiple cities, like the Washington, DC–Baltimore metropolitan area.

Disease Spread Dynamics

Within each city and rural area, the progression of the flu is described by an SEIR model.[15] Exposed or infected individuals exhibit no symptoms, are not yet capable of infecting others, but will eventually become infectious. Individuals who are infected and infectious are distinguished according to the number of days that have elapsed since they were first exposed (τ) or since they became infectious (θ), and by whether they are symptomatic (s=1) or asymptomatic (s=0). Thus $E(\tau, t)$ denotes the number of people who are exposed on day t who became infected on day t-τ.$I(s,\theta,t)$ denotes the number of people who are infectious on day t who first became infected on day t-θ. The maximum length of the latent period is $\bar{\tau}$ days; the maximum length of time between first being infected and recovering or dying is $\bar{\theta}$ days.

The stocks of susceptible, exposed, infectious, recovered, and dead persons in city i on day t is given by the right-hand side of (12), which must equal the city population, N. (From this point forward we suppress the city i index for notational simplicity.)

$$N(t) = S(t) + \sum_{\tau=0}^{\bar{\tau}} E(\tau, t) + \sum_{\tau=0}^{\bar{\theta}} \sum_{s=0}^{1} I(s, \theta, t) + R(t) + D(t). \quad (12)$$

The key equations of the model describe the change in the stock of infectious, exposed, and susceptible people from one day to the next. Define $\gamma(\tau)$ as the probability that an exposed case transitions to being infectious and $\sigma(\theta)$ as the probability that an infectious person recovers. We assume for simplicity (following Longini et al. 2005) that 80 percent of exposed cases transition to infectious after one day and 20 percent after two days; they will become symptomatic infectious with a probability of 0.67 and asymptomatic infectious with a probability of 0.33. Infectious cases either recover or die three to six days after they become infectious, with 30 percent recovering after three days, 40 percent after four days, 20 percent after five days, and 10 percent after six days. This means $\bar{\tau} = 1$ and $\bar{\theta} = 5$.

[15] To simplify notation, we present the model for a single age category. In the GEM, all population variables vary by age category.

Thus:

$$\gamma(\theta) = \begin{cases} 0.8 \; if \; \tau = 0, \\ 1 \; if \; \tau = 1, \\ 0 \; otherwise. \end{cases} \tag{13}$$

$$\sigma(\theta) = \begin{cases} 0.3 \; if \; \theta = 2, \\ \frac{0.4}{0.7} if \; \theta = 3, \\ \frac{0.2}{0.3} if \; \theta = 4, \\ 1 \; if \; \theta = 5, \\ 0 \; otherwise. \end{cases} \tag{14}$$

Ignoring travel between cities, the number of infectious cases in city *i* evolves according to

$$I(1,\theta,t+1) = \begin{cases} (0.67)\sum_{\tau=0}^{\bar{\tau}} [E(\tau,t)\gamma(\tau)], for \; \theta = 0 \\ I(1,\theta-1,t)\big(1-\sigma(\theta-1)\big), for \; \theta = 1, \ldots \bar{\theta} \end{cases} \tag{15}$$

$$I(0,\theta,t+1) = \begin{cases} (0.33)\sum_{\tau=0}^{\bar{\tau}} [E(\tau,t)\gamma(\tau)], for \; \theta = 0 \\ I(0,\theta-1,t)\big(1-\sigma(\theta-1)\big), for \; \theta = 1, \ldots \bar{\theta} \end{cases} \tag{16}$$

for symptomatic and asymptomatic cases, respectively.

We assume 97.5 percent of symptomatic cases recover and 2.5 percent die, while 98.75 percent of asymptomatic cases recover and 1.25 percent die.[16] This implies that the number of recovered cases evolves according to

$$R(t+1) = R(t) + \sum_{\tau=0}^{\bar{\tau}} \sum_{s=0}^{1} \big[\big(0.975s + 0.9875(1-s)\big) I(s,\tau,t)\sigma(\tau)\big], \tag{17}$$

[16] This assumption is highly speculative, but gross attack rates would be similar so long as case fatality rates remained small.

while the number of dead cases evolves according to

$$(18) \quad D(t+1) = D(t) + \sum_{\theta=0}^{\bar{\theta}} \sum_{s=0}^{1} \big(0.025s + 0.0125(1-s)\big) I(s,\theta,t)\sigma(\theta).$$

A new exposed case is the result of contact between a susceptible person and an infectious person. Groups within a city mix uniformly, so the average number of exposed cases caused by one infectious person on day *t* is proportional to the number of susceptible individuals in the city and the daily infectious contact rate λ, which is equivalent to the parameterβ in the simple SIR models. (In reality, λ varies by age group, as described below.) Define *r(s)* as the relative infectiousness of symptomatic state *s* and *f(t)* as a seasonality factor. Based on Ferguson et al. (2005) and Longini et al. (2005), we assume *r(s)* = 1 if a person is symptomatic (s=1) and *r(s)* = 0.5 if a person is asymptomatic (s=0). We discuss seasonality below.

The number of exposed cases in city *I* evolves according to the following:

$$(19) \quad E(\tau, t+1) = \begin{cases} \lambda f(t) \dfrac{S(t)}{N(t) - D(t)} \displaystyle\sum_{\theta=0}^{\bar{\theta}} \sum_{s=0}^{1} [I(s,\theta,t) \cdot r(s))], for\ \tau = 0, \\ E(\tau - 1, t)\big(1 - \gamma(\tau - 1)\big), for\ \tau = 1, \ldots \bar{\tau}. \end{cases}$$

The preceding describes an S-E-I-R model with no air travel. In the model, the evolution of S(.), E(.), and I(.) are determined by both these equations and by travel between cities. There is a fixed number of seats available X_{ij} for travelers moving between two cities *i* and *j* determined either from the airline data (for city to city travel) or 1% of the urban population (for city to rural travel). Each susceptible, exposed, asymptomatic infectious, and recovered individual has a probability of travel calculated by the following:

$$(20) \quad pTravel(i,j) = \frac{X(i,j)}{N(i, t=0)}.$$

The expected number of travelers of a particular subgroup who move from *i* to *j* on a particular day is the product of the number of persons in the subgroup multiplied by $pTravel(i,j)$.[17]

[17]As a result, the expected number of travelers moving from *i* to *j* will be less than X_{ij} as long as any cases are symptomatic infectious or dead because these people do not travel.

Influenza Transmissibility

The GEM differs from the model described above by distinguishing four population age groups: 0–4 (group 0), 5–14(group 1), 15–64(group 2), and 65 or older (group 3). Members of all age groups and disease states are assumed to mix randomly within a city, although contact rates differ among age groups.

The key parameter determining the spread of the pandemic is λ_{ijk}, the (city-specific) daily infectious contact rate. This parameter defines the number of new exposed cases of age *k* generated by an infectious individual of age *j* in city *i*. The rate is determined by two components: the number of people contacted daily by an infectious individual and the probability of transmission given contact, *P(T|C)*. The contact rates are described by a matrix C_i whose elements C_{ijk} describe the number of contacts per day that a person of age *j* has with people of age *k* in city *i*. C_{ijk} is modified by a factor d_i to reflect the higher number of contacts experienced by all age groups in more densely populated cities:

(21) $\lambda_{ijk} = d_i \cdot \sum_{a=0}^{3} C_{iak} \cdot P(T|C)$, *j=0, 1, 2, 3.*

The elements of the contact rate matrix are derived from data on contact rates from Mossong et al. (2008), as described in Appendix B. In general each age group is more likely to contact people of their own age than other ages. School-age children have the highest contact rates, while people 65 and older have the lowest contact rates. This pattern influences the incidence of influenza cases across age groups. However, we assume that the total absolute number of contacts per day does not vary across cities based on age structure, implying that countries with younger populations will not have higher attack rates. We adjust contact rates across cities by urban density and across rural areas by their urbanization rate (see Appendix B for details).This increases contact rates (and thus attack rates) for cities in lower-income countries, which tend to have high-density cities, and rural areas in higher-income countries, which have higher urbanization rates.

The probability of transmission given contact follows particular biological characteristics of the influenza strain, does not vary with age, and determines the reproductive rate of the

virus.[18] The values chosen ($P(T|C)$ = 0.045,0.0533 and 0.06) correspond to global attack rates in the absence of treatment of 28, 45, and 53 percent.

Seasonality of the Flu

The seasonality factor $f(t)$ (see equation (19)) reflects the fact that influenza infection rates are much higher in winter months than in summer in temperate zones but that seasonal variation is dampened in tropical or arctic conditions. The mechanism for this lack of variation is not well understood. Lowen et al. (2008) have confirmed in controlled environment tests that the influenza virus is more fragile in warmer than colder temperatures and has lower transmission at higher humidity, perhaps because virus particles are coated in moisture and less able to infect new hosts.[19] Behavioral mechanisms suggest that in winter months (particularly in higher latitudes) people spend more time indoors in close proximity, and so have higher contact rates and thus a higher daily infectious contact rate.

All global influenza models attempt to account for seasonal variation; most models use some kind of step function as an approximation, where seasonality is constant for tropical regions and is lower during summer months in the temperate zones (e.g.,Rvachev and Longini 1985; Grais et al. 2003; Cooper et al.2006; Colizza et al. 2007). Because of the generally poor understanding of seasonal factors, we follow the general practice of the global influenza literature.

The GEM incorporates seasonality through a seasonal infectiousness factor that varies by day and latitude.[20] We assume infectiousness is constant in the tropics; elsewhere, it follows a sinusoidal pattern where it is at maximum value on January 1 in the Northern Hemisphere (July 1 in the Southern Hemisphere) and minimum value on July 1 in the Northern Hemisphere (January 1 in the Southern Hemisphere), as illustrated in Figure 7. This implies that disease transmission, and thus attack rates, are highest in the tropics and are lower the further a country is from the equator (up to the polar boundary).

---

[18] A simple model extension would be to allow $P(T|C)$ to vary by age groups so that, for example, a child or elderly individual might be more likely to contract the virus than a healthy adult. But reliable data on relative vulnerabilities are not available, and such a model change would be functionally identical to changing the elements of the contact matrix, since what we care about is the total effect of contact and infection.

[19] Behavioral mechanisms suggest that in winter months (particularly in higher latitudes), people spend more time indoors in close proximity and so have higher contact rates and thus a higher infectiousness. Here, we focus on adjusting infectiousness rates rather than contact rates.

[20] We use the mean latitude of the modeled cities in a region for the latitude of that region's rural area.

Figure 7 shows the seasonality factors for the tropical city of Singapore, the temperate city New York in the Northern Hemisphere, and the temperate city of Sydney in the Southern Hemisphere with a pandemic starting on January 1. Note that New York is further north than Sydney is south, so New York's seasonality factor has higher amplitude and a lower mean value. These imply that all else equal, Singapore will have a higher attack rate than Sydney, which will have a higher attack rate than New York.

Unfortunately, there is little empirical data on which to base this assumption. Cooper et al. (2006) compare a constant-maximum infectiousness model (as latitude changes) to a constant-mean model and finds the former performs slightly better when trying to fit observed attack rates from the 1968/69 influenza pandemic. They also find that a sine wave functional form fits the data better than a square wave or no seasonal variation. These results support the assumptions made in this paper. In contrast, evidence from Shaman and Kahn (2009) suggests that influenza seasonality is driven largely by absolute humidity. Absolute humidity is always high in the tropics but is high in temperate zones only in summer, which suggests that a constant minimum model may be more accurate than a constant-maximum model. But, further evidence is lacking. This is unfortunate, since this seasonality assumption has a huge impact on pandemic dynamics and on the incidence of attack rates across countries.

Figure 7. Seasonality Factor for Selected Cities

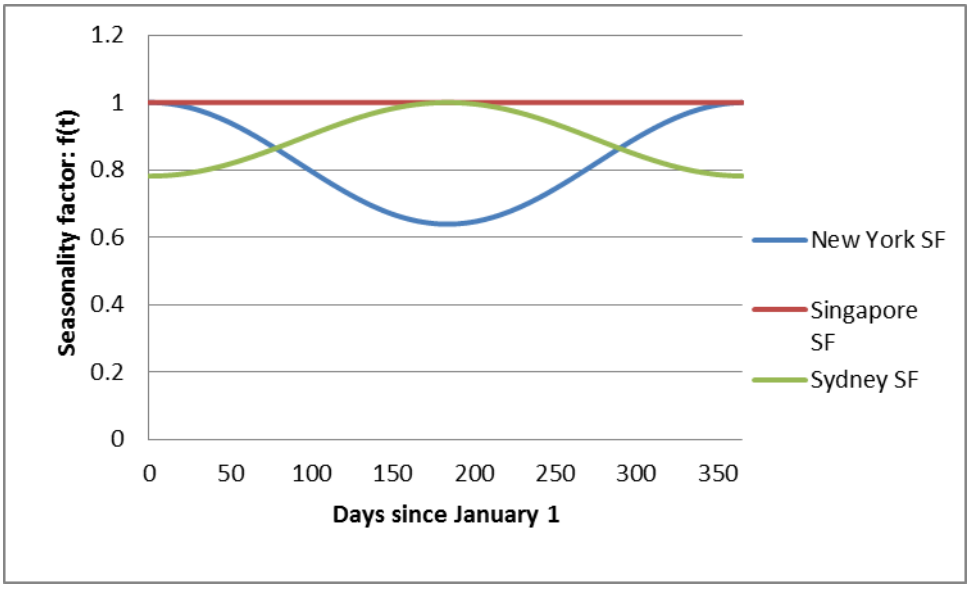

Model Runs with No Treatment

In our base case, we assume the pandemic begins with 100 infectious cases on January 1 in Jakarta, Indonesia. We select Jakarta because it is a city in a large, low-income country that is well-connected to the international airline network and because a pandemic influenza strain that mutates from an avian or other animal strain is likely to occur in a country where people in rural areas live in close proximity to livestock, as they do in Indonesia. Although the model is run at a daily time-step for a three-year period, we focus on the attack rate after one year. This represents the time that it might take to develop, test, and mass produce a vaccine specific to the pandemic strain.

Every region in the model is assigned to an income group — low, lower middle, upper middle, or high (see Table 2 and Appendix C)—based on International Monetary Fund per capita 2008 GDP data. Low-income countries are those with per capita GDP less than USD$3,000;lower-middle-income countries up to $10,000; upper-middle countries, up to $20,000; and high-income countries, more than $20,000. Regions with multiple countries are classified by the median-income country. We use this income group classification rather than the World Bank's because we feel it is more representative of public health infrastructure and therefore useful for defining antiviral scenarios where the size of the stockpile differs depending on the purchasing power and public health system of the country in question.

Table 2. Distribution of Population by Income and Age Group

Income group			Population		
	Total	Ages 0–4	Ages 5–14	Ages 15–64	Ages 65+
Low	2,839,998,810	352,152,534	647,655,318	1,716,456,775	123,734,183
Lower middle	2,156,161,724	165,680,600	361,106,698	1,473,910,542	155,463,884
Upper middle	515,385,759	36,733,909	79,318,428	349,989,705	49,343,717
High	914,861,650	53,135,831	109,758,061	610,674,513	141,293,245
TOTAL	6,426,407,943	607,702,874	1,197,838,505	4,151,031,535	469,835,029

The baseline scenario assumes a moderate *P(T|C)* value of 0.05333, which results in a world global attack rate of 45 percent at the end of one year in the no-intervention case. In a global model, the number of people infected by one individual in an otherwise susceptible

population (R_0) varies across cities and regions, making direct comparisons difficult between $P(T|C)$ and R_0 in the models of sections 3 and 4.

Figure 8 shows how the pandemic progresses in the no-treatment scenario. Panel A shows the number of new infectious cases each day, by income group, for the duration of the pandemic, while panel B shows the daily number of new infectious cases the first year of the pandemic. Panel C shows attack rates at the end of the first year by county income group, while panel D shows first-year attack rates by age group. The pandemic's worldwide peak occurs around day 200, with a secondary wave of infection occurring in the second year, peaking around day 400 as the Northern Hemisphere re-enters winter. Most cases occur within a year; the one-year attack rate is 44.6 percent, rising to 49.9 percent after three years.

Attack rates are higher in low-income countries and lower in high-income countries, primarily due to the positive correlation of income and latitude. High-income, higher latitude countries have a larger increase in attack rate after the first year because they are more affected by seasonality. School-age children suffer higher attack rates than other age groups because of their high contact rates; persons 65 and older suffer lower attack rates because of their relatively low contact rates.[21]

Figure 8. Pandemic Time Path, No Treatment

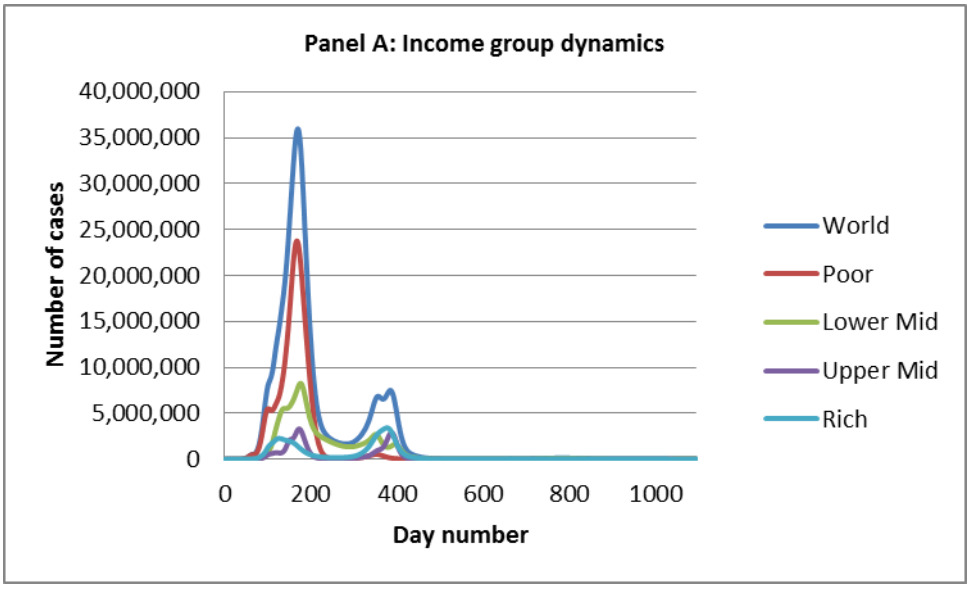

[21] These patterns match data on attack rates in the United States during the three major pandemics of the 20th century (Glezen 1996).

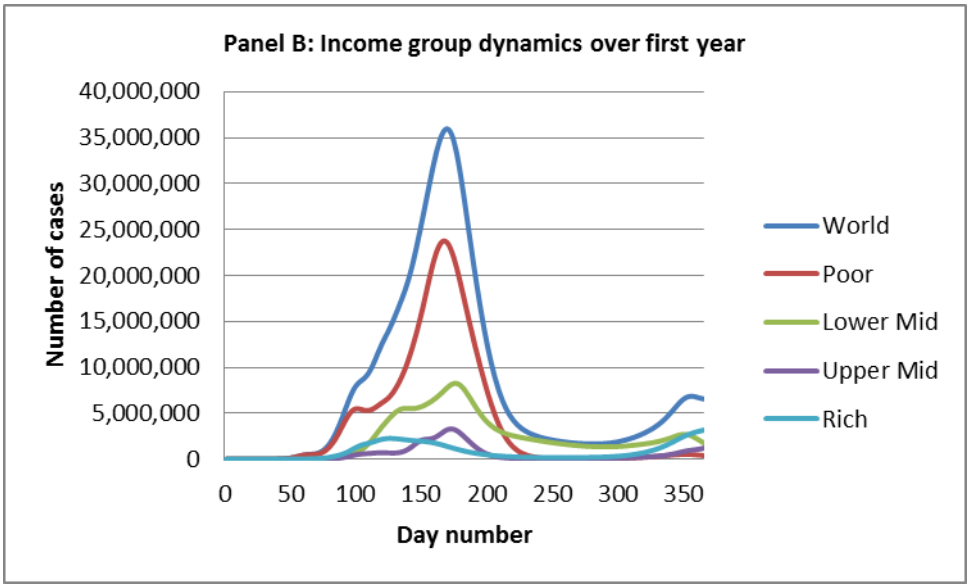
Panel B: Income group dynamics over first year
40,000,000
35,000,000
30,000,000
25,000,000
20,000,000
15,000,000
10,000,000
5,000,000
0
Number of cases
0
50
100
150
200
250
300
350
Day number
World
Poor
Lower Mid
Upper Mid
Rich

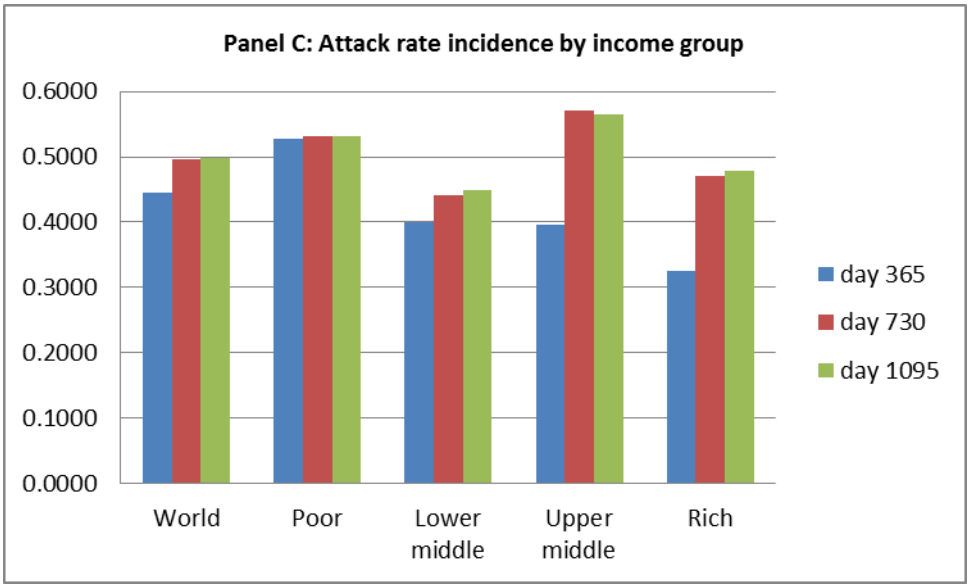
Panel C: Attack rate incidence by income group
0.6000
0.5000
0.4000
0.3000
0.2000
0.1000
0.0000
World
Poor
Lower middle
Upper middle
Rich
day 365
day 730
day 1095

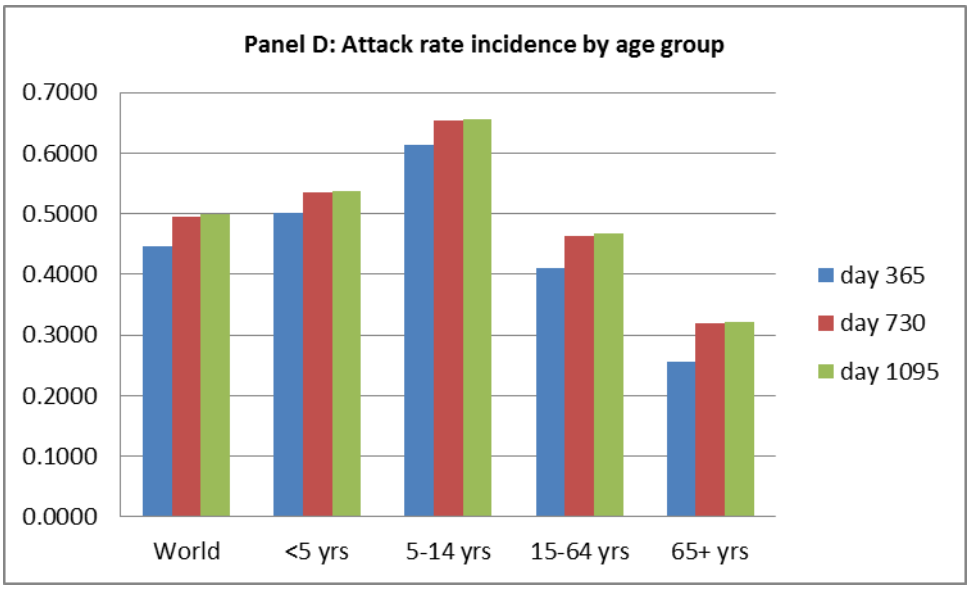

Attack rates and their incidence by county group are highly sensitive to the pandemic start date due to the interaction of start date and seasonality assumptions. Figure 9 compares the pandemic time path and attack rates for a January 1 start date to a July 1 start date. The secondary infection wave occurs in the January 1 scenario only because the pandemic is choked off prematurely by rising temperatures as the Northern Hemisphere enters summer, while in the July1 scenario, Northern Hemisphere temperatures are cooling as the disease spreads, so the pandemic reaches a single, more intense peak. The increase in the one-year attack rate is largest for high-income countries since their high latitude means this choking off effect is stronger.

Figure 9. Sensitivity of Pandemic Dynamics to Start Date

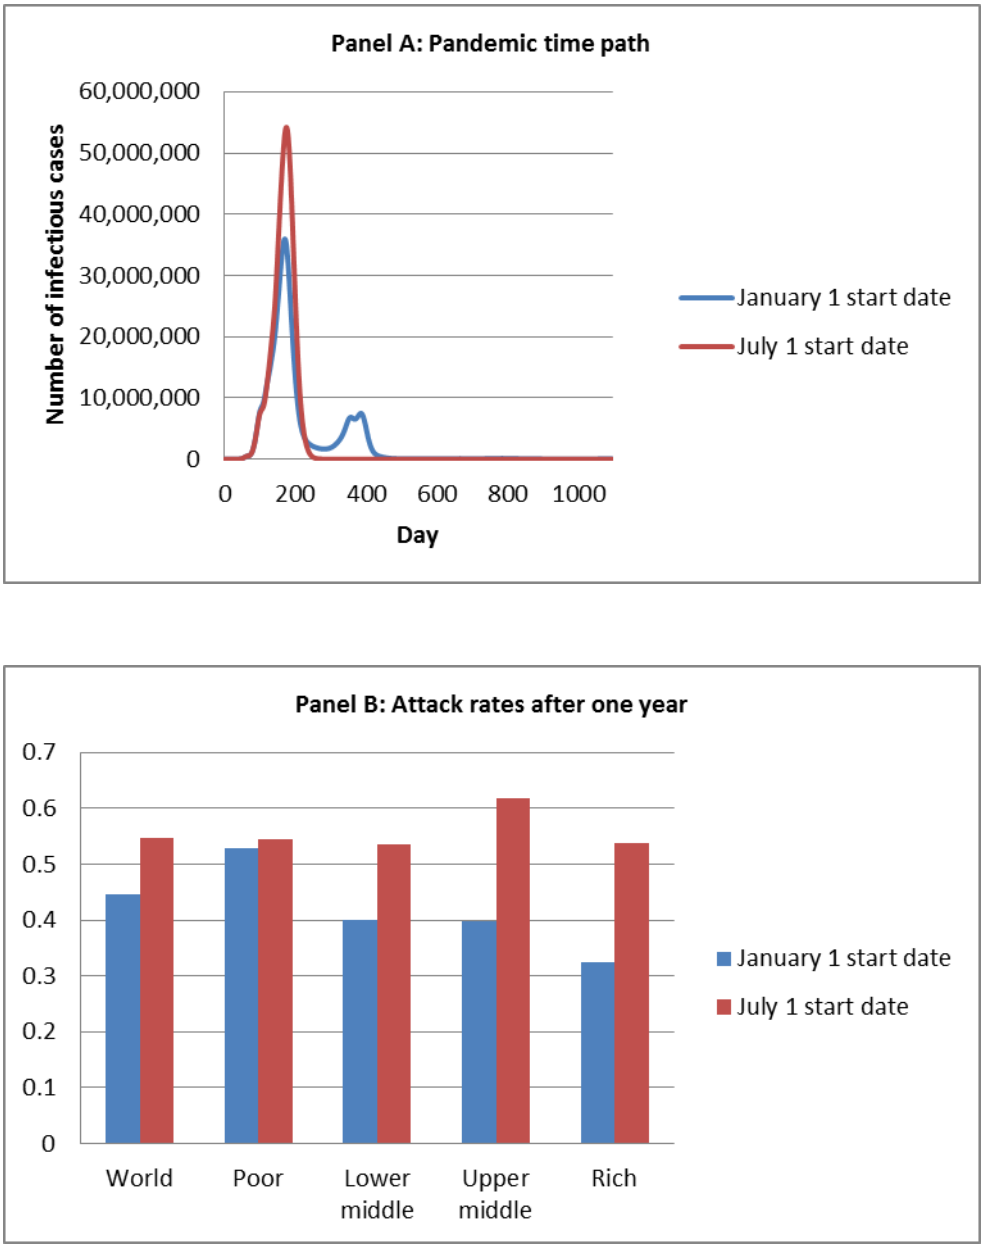

The incidence of attack rates across income groups is driven primarily by the combination of the seasonality assumptions and the correlation between income and latitude. Higher-income countries lie farther from the equator, where they have lower average seasonality factors and attack rates. Figure 10 displays the distribution of attack rates (after three years) by latitude across the GEM's 106 regions. Panel A shows the distribution of regions across latitude by income group and demonstrates the stylized fact that higher-income countries lie farther from

the equator. Panel B shows attack rates under a January 1 start date, while panel C shows attack rates under a July 1 start date. The decrease in attack rates is larger for the Northern Hemisphere in panel B and the Southern Hemisphere in panel C because of the timing of the pandemic peak relative to summer, when the seasonality factor is at a minimum. Variation in attack rates not due to latitude can be explained by the structure of the international airline network as well as variation in population density and urbanization rates across countries. The upward slope above the mid-30s latitude in Panel B is an example.

Figure 10. Distribution of Attack Rate by Latitude

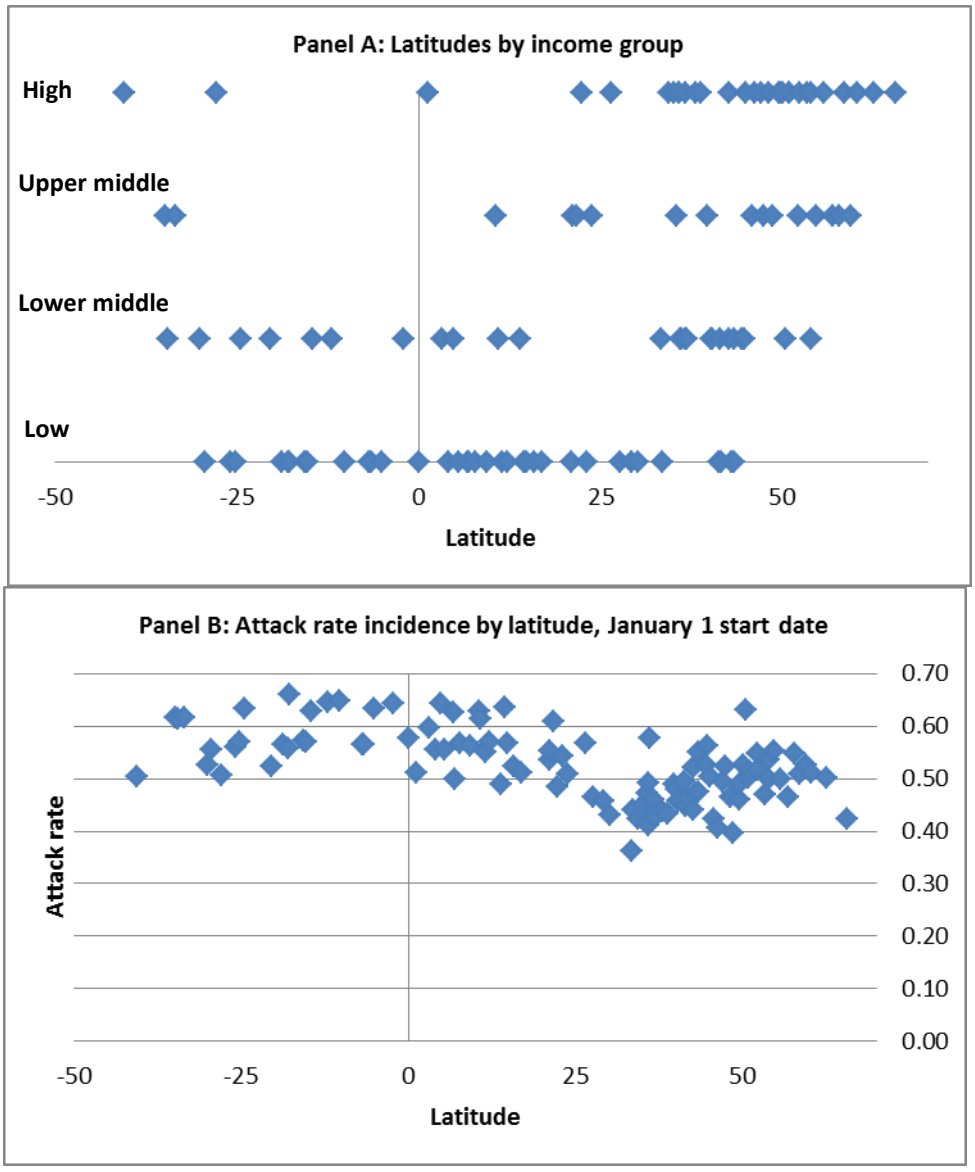

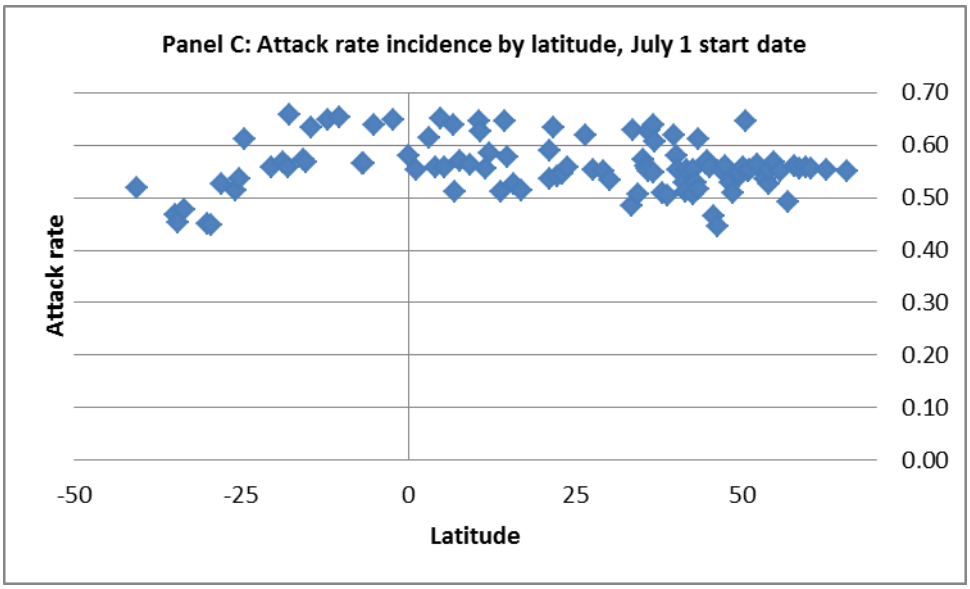

A higher value of *P(T|C)* results in a faster pandemic and a higher attack rate than a lower *P(T|C)* value while retaining similar qualitative properties (see Figure 11). In a severe pandemic with *P(T|C)* = 0.06, the global one-year attack rate rises to 53.3 percent, while in a minor pandemic with *P(T|C)* = 0.045, the global one-year attack rate falls to 28.7 percent.

Figure 11. Sensitivity of Pandemic Dynamics to Infectiousness, *P(T|C)*

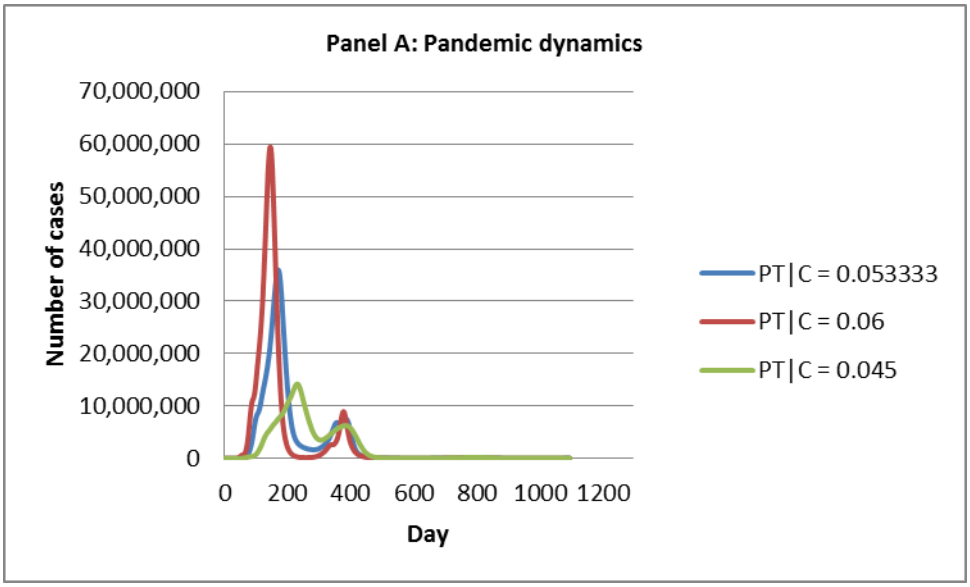

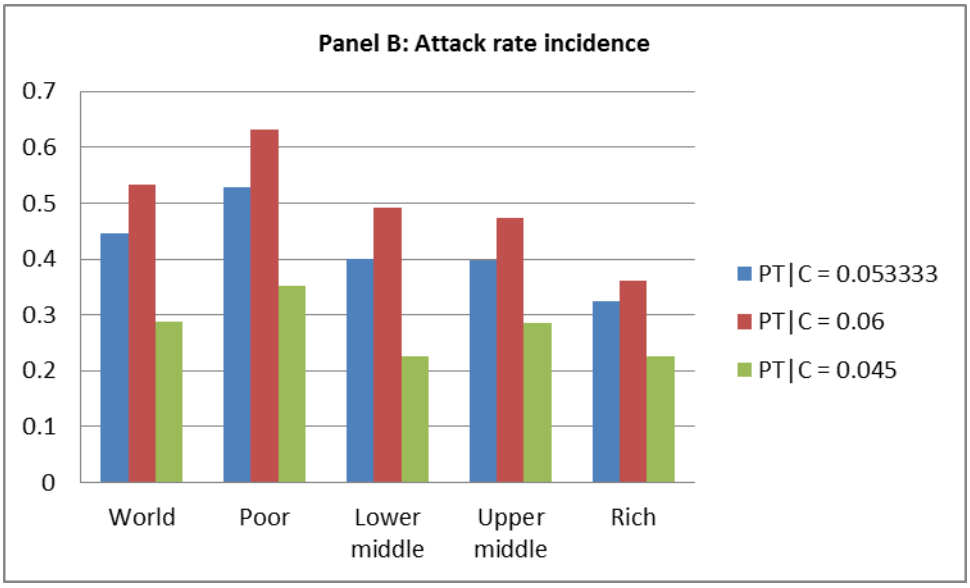

Results are generally insensitive to the travel rate. Figure 12 shows that one-year attack rates for wealthier countries vary slightly with a doubling or halving of travel rates and the impact on terminal attack rates after three years is minimal. A higher travel rate has two effects: a slightly faster pandemic spread (which might be expected to lead to higher attack rates) and a shift in the season in which the pandemic peak occurs (which may increase or decrease attack rates depending on the season). With a January 1 start date (which peaks in Northern Hemisphere summer) the second effect appears to dominate: attack rates are marginally lower with higher travel rates. The apparent insensitivity of attack rates to travel rates implies that airline travel restrictions would be relatively ineffective (Epstein et al. 2007).

Figure 12. Sensitivity of Pandemic Dynamics to Travel Rates

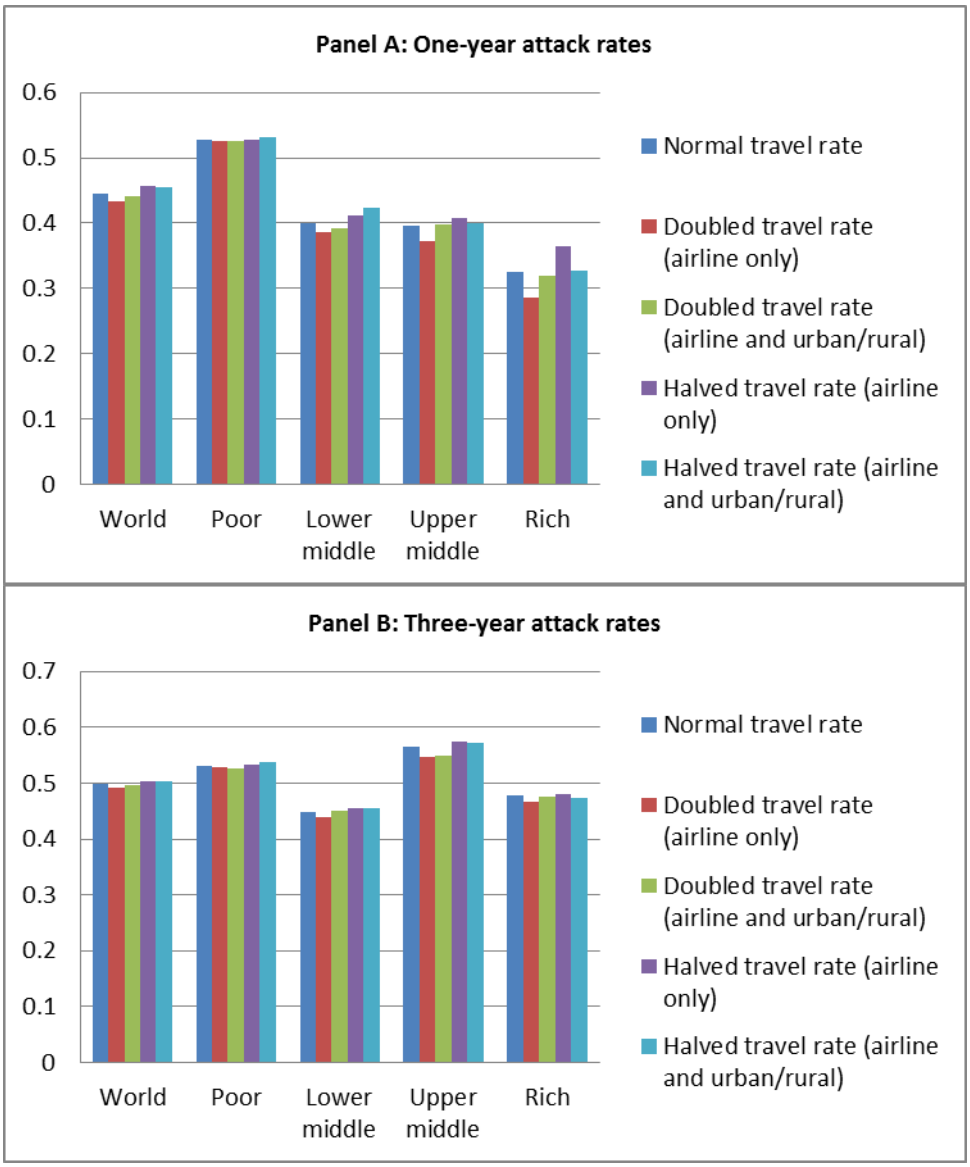

6. Policy Interventions in the Global Epidemiological Model

We use the GEM to examine the impacts of various antiviral stockpile scenarios. In all scenarios, high-income and upper-middle-income countries have stockpiles. Whether stockpiles are distributed to lower-middle-income or low-income countries varies among scenarios. We examine the nature of externalities associated with stockpile use: how does the use of antivirals by one group of countries affect the attack rate in other groups of countries?

Details of Antivirus Treatment

Antivirals held by region *r* are used to treat people in all cities and the rural area within that region. We assume that antiviral stockpile distribution begins in city *i* (or in the rural area) only when 1,000 people have become infectious and then continues as needed until no antiviral remains in the region or every infectious person has been treated.[22] Antivirals are distributed to a proportion p^* of symptomatic infectious people.[23] In most runs, $p^* = 0.5$, implying that one-third of total new infectious (including asymptomatic) cases are treated each day, provided that the stockpile has not run out. Treatment occurs after an individual has been infectious for one day.

Following Longini et al. (2005), we assume that treating infectious cases with antivirus has two effects. First, it reduces the probability of transmission given contact by a multiplicative factor $(1 - e)$, where e is the efficacy of the antivirus. Second, it reduces the average duration spent infectious by one day because it uniformly shifts the probability distribution downward. As a result, an infectious individual treated with antivirals will recover two to five days after becoming infectious, with 30 percent recovering after two days, 40 percent after three days, 20 percent after four days, and 10 percent after five days. In most runs e=0.6.

We assume antiviral doses are never used to treat non-infectious cases. While unrealistic, this assumption is relatively harmless because we could increase the stockpile size in every city by some factor and obtain the same qualitative results.

Choice of Antiviral Scenarios

We focus on four antiviral stockpile scenarios based on income group, where the size of the stockpile is a percentage of the population of each group (low/lower middle/upper middle/high).[24] The four scenarios are (A) 0/0/5/10; (B) 0/1/5/10; (C) 1/1/5/10; and (D) 0/1/5/10,

---

[22] This assumption is designed to model the difficulty in detecting the pandemic strain of influenza, particularly against the background noise of regular seasonal influenza cases. Sensitivity tests show that final attack rates are not significantly affected by changing the number of cases needed before distribution begins, except in the special case of low $P(T|C)$, where early detection will significantly reduce the number of cases, particularly in the country where the outbreak originates.

[23] We do not consider prophylactic use of antivirals. Although a susceptible person treated with an antiviral has a lower probability of being infected if contacted, the effect of being treated wears off rapidly after treatment ceases. This implies that doses could be more effectively used for treating actual infectious cases.

[24] Although all regions in a particular income group receive the same stockpile, as a percent of their population, it is possible that the stockpile may be exhausted in some regions but not in others. This complicates interpretation of a change in stockpile size across an income group because such a change may lead to additional people being treated in some regions but not in others.

plus a 4.2 percent stockpile for the low-income outbreak country, Indonesia. In all scenarios, high-income regions have a combined stockpile of 91.5 million doses, and upper-middle-income regions have a stockpile of 25.8 million doses (10 and 5 percent of their respective populations).[25] The difference between scenario A and scenario B could be interpreted as a gift from wealthier countries of 21.5 million doses to lower-middle-income countries (equal to 1 percent of their population), while the movement from B to C is equivalent to high-income countries paying for 28.4 million doses in low-income countries (equal to 1 percent of their population).Scenario D mirrors scenario B except that Indonesia has a stockpile of 9.2 million doses, which equals 4.2 percent of its population.

The impact of each scenario on world attack rates and attack rates across country groups at the end of one year depends on three factors: the infectiousness of the flu ($P(T|C)$), the outbreak date, and the effectiveness of antivirals (e).We therefore present sensitivity analyses for all scenarios while describing their impact.

Effect of Antivirals on World Attack Rates

Treatment with antivirus is highly effective in reducing pandemic attack rates. Table 3 shows the effect of antiviral treatment under scenarios A–C compared to a scenario with no treatment. The table also shows the impact of varying $P(T|C)$, the start date of the pandemic, and antiviral efficacy. Tables 4–6 show the impact of varying these parameters on the cumulative number of influenza cases at the end of one year. In our standard case (medium infectiousness, January 1 start date, and reduction in infectiousness of 60 percent), Scenario A reduces the attack rate by 5.2 percentage points (334 million cases); scenario B reduces the attack rate by 13.6 percentage points (877 million cases); and scenario C reduces the attack rate by 19.3 percentage points (1,237 million cases), all relative to the baseline without treatment. Scenario A leads to an average of 2.85 fewer cases per dose; scenario B, 7.16 fewer cases per dose; and scenario C, 8.60 fewer cases per dose.

[25] However, note that the entire 91.5 million doses are not always consumed, or are not always consumed within the first year, as some regions will not exhaust a stockpile equal to 10 percent of their population size. Since one-third of infectious cases are treated, a region will not exhaust a 10 percent stockpile any time its attack rate is less than 0.3.

Table 3. Effect of Antivirals (AV) on Global One-Year Attack Rate

Scenario ID	AV stockpile size (% population, by income group)	*P(T|C)*	Start date	AV efficacy	Baseline attack rate	Attack rate with AV	Difference
A1	0/0/5/10	Low	1-Jan	0.6	0.287	0.235	0.052
A2	**0/0/5/10**	**Medium**	**1-Jan**	**0.6**	**0.446**	**0.394**	**0.052**
A3	0/0/5/10	Medium	1-Jul	0.6	0.547	0.456	0.091
A4	0/0/5/10	Medium	1-Jan	0.5	0.446	0.400	0.047
A5	0/0/5/10	High	1-Jan	0.6	0.533	0.480	0.053
B1	0/1/5/10	Low	1-Jan	0.6	0.235	0.154	0.081
B2	**0/1/5/10**	**Medium**	**1-Jan**	**0.6**	**0.446**	**0.310**	**0.136**
B3	0/1/5/10	Medium	1-Jul	0.6	0.547	0.385	0.162
B4	0/1/5/10	Medium	1-Jan	0.5	0.446	0.325	0.121
B5	0/1/5/10	High	1-Jan	0.6	0.533	0.499	0.034
C1	1/1/5/10	Low	1-Jan	0.6	0.235	0.000	0.234
C2	**1/1/5/10**	**Medium**	**1-Jan**	**0.6**	**0.446**	**0.254**	**0.193**
C3	1/1/5/10	Medium	1-Jul	0.6	0.547	0.301	0.246
C4	1/1/5/10	Medium	1-Jan	0.5	0.446	0.309	0.137
C5	1/1/5/10	High	1-Jan	0.6	0.533	0.461	0.072

Note: "Standard" assumption scenarios are in bold. AV stockpile sizes refer to the following income regions: low/lower middle/upper middle/high. Low P(T|C) = 0.045, medium P(T|C) = 0.05333, and high P(T|C) = 0.06.

Table 4. Sensitivity of Reduction in Global Influenza Cases from Antivirals (AV) to Infectiousness

	P(T|C) = 0.045			*P(T|C)* = 0.05333			*P(T|C)* = 0.06		
Scenario	Δ attack rate	Δ cases (billions)	Cases reduced per AV dose	Δ attack rate	Δ cases (billions)	Cases reduced per AV dose	Δ attack rate	Δ cases (billions)	Cases reduced per AV dose
Baseline (no AV)	0	0		0	0		0	0	
Scenario A	0.052	0.333	2.836	0.052	0.334	2.848	0.053	0.341	2.906
Scenario B	0.132	0.850	6.946	0.136	0.877	7.163	0.034	0.219	1.793
Scenario C	0.286	1.839	12.773	0.193	1.237	8.594	0.072	0.466	3.235

Note: Change in attack rate and number of cases refer to the reduction relative to the baseline without AV. All figures are for the point one year after the pandemic commences.

Table 5. Sensitivity of Reduction in Global Influenza Cases from Antivirals (AV) to AV Efficacy

	e = 0.6			*e* = 0.5		
Scenario	Δ attack rate	Δ cases (billions)	Cases reduced per AV dose	Δ attack rate	Δ cases (billions)	Cases reduced per AV dose
Baseline (no AV)	0	0		0	0	
Scenario A	0.052	0.334	2.848	0.047	0.299	2.550
Scenario B	0.136	0.877	7.163	0.121	0.779	6.360
Scenario C	0.193	1.237	8.594	0.137	0.879	6.104

Note: Change in attack rate and number of cases refer to the reduction relative to the baseline without AV. All figures are for the point one year after the pandemic commences.

Table 6.Sensitivity of Reduction in Global Influenza Cases from Antivirals (AV) to Start Date

	January 1 start			July 1 start		
Scenario	Δ attack rate	Δ cases (billions)	Cases reduced per AV dose	Δ attack rate	Δ cases (billions)	Cases reduced per AV dose
Baseline (no AV)	0	0		0	0	
Scenario A	0.052	0.334	2.848	0.091	0.585	4.993
Scenario B	0.136	0.877	7.163	0.162	1.041	8.508
Scenario C	0.193	1.237	8.594	0.246	1.581	10.984

Note: Change in attack rate and number of cases refer to the reduction relative to the baseline without AV. All figures are for the point one year after the pandemic commences.

Increasing the infectious of the flu ($P(T|C) = 0.06$) reduces the effectiveness of antivirals, while a lower value of $P(T|C)$ makes containment of the flu possible under scenario C. Scenario C allocates antivirals to all low-income countries, including the outbreak source (Indonesia), and the low $P(T|C)$ value implies that treatment in Indonesia is able to reduce the R_0 value below 1, preventing a pandemic. But containment is not possible in other scenarios. Reducing the effectiveness of the antiviral dose from e=0.6 to e=0.5 lowers the number of cases reduced per dose of antiviral from 2.85 to 2.55 in scenario A, from 7.16 to 6.36 in scenario B, and from 8.60 to 6.10 in scenario C.

The start date of the pandemic also affects the impact of antivirals on the global attack rate. Figure 13 shows how antiviral treatment delays the pandemic, shifting the distribution to the right. With a January 1 start date, delay blunts the effectiveness of antiviral in reducing the attack rate because it pushes the pandemic's peak toward Northern Hemisphere winter, when the

seasonality factor is near maximum. With a July 1 start date, delay augments the effectiveness of antiviral in reducing the attack rate because it pushes the pandemic peak away from Northern Hemisphere winter.

Figure 13. Pandemic Time Path with Antiviral Treatment

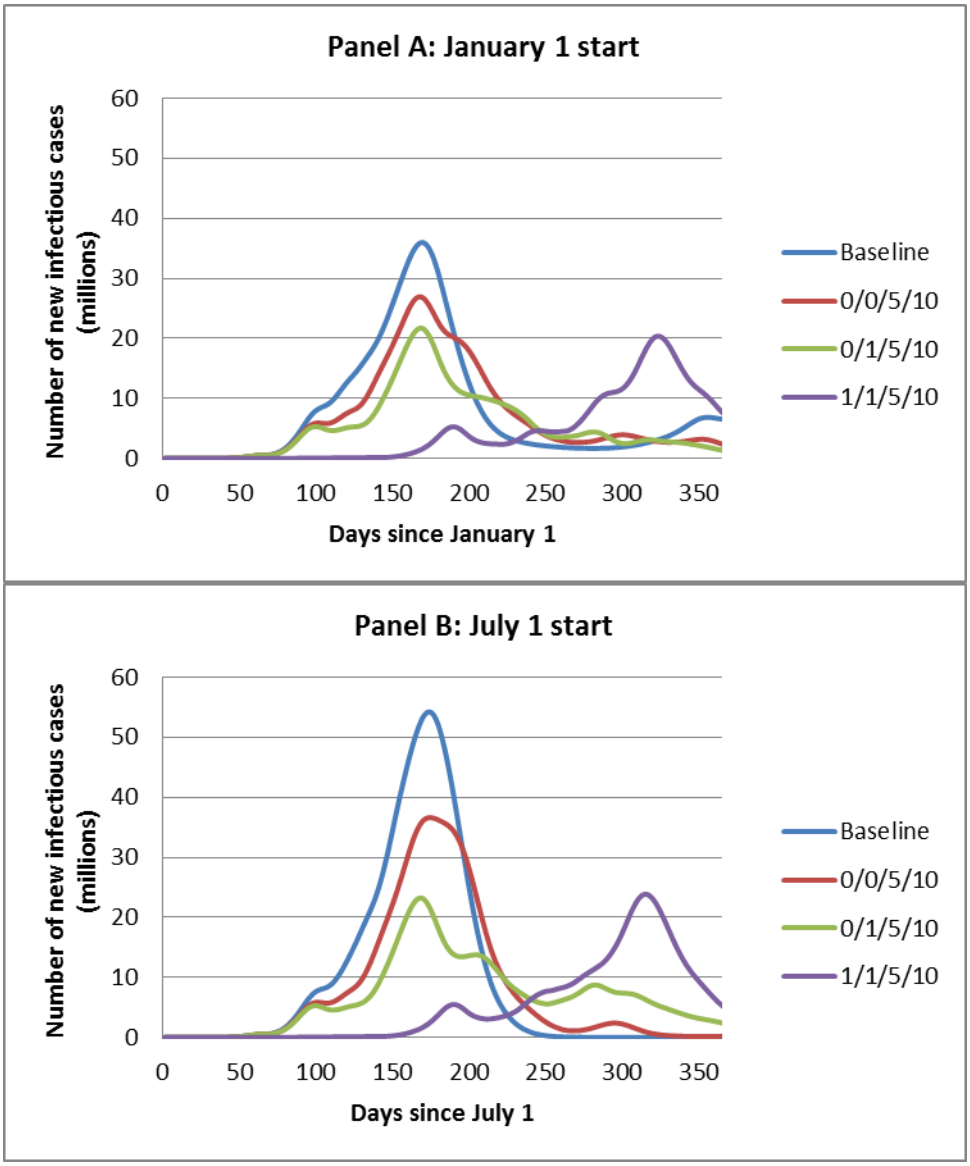

Effect of Antiviral Usage in One Region on Other Regions

In one and two-city SIR models, most of the gains from antiviral treatment accrued to the region in which the treatment was given, but some positive effects spilled over to other regions.

These results also hold in the full GEM. Table 7shows the reduction in attack rates across income groups for the base case assumptions, as well as variation in transmissibility of the flu and start date. Moving from no antiviral to the 0/0/5/10 scenario in the base case, high-income and upper-middle-income regions experience a large reduction in attack rate, while lower-middle-income and low-income regions see negligible impacts.[26] Adding a stockpile in lower-middle-income regions significantly reduces their attack rates while causing a small reduction in high-income country attack rates and a larger reduction in attack rates in upper-middle-income countries.[27] Moving to 1/1/5/10 adds an antiviral stockpile in low-income countries, leading to a large reduction in their attack rates and moderate benefits elsewhere.

Table 7. Impact of Antivirals (AV) on One-Year Attack Rate, by Income Group

AV scenario	#AV doses	World	Low	Lower middle	Upper middle	High
January 1 start, medium *P(T|C)*						
Baseline (no AV)	0.0	0.446	0.528	0.401	0.397	0.325
Scenario A	117.3	0.394	0.529	0.403	0.210	0.059
Scenario B	138.8	0.310	0.525	0.169	0.166	0.052
Scenario C	167.2	0.254	0.458	0.121	0.051	0.045
January 1 start, low *P(T|C)*						
Baseline (no AV)	0.0	0.287	0.353	0.226	0.285	0.225
Scenario A	117.3	0.235	0.351	0.223	0.010	0.027
Scenario B	138.8	0.154	0.343	0.003	0.003	0.010
Scenario C	167.2	0.000	0.001	0.000	0.000	0.000
January 1 start, high *P(T|C)*						
Baseline (no AV)	0.0	0.533	0.631	0.492	0.473	0.362
Scenario A	117.3	0.480	0.630	0.487	0.299	0.102
Scenario B	138.8	0.499	0.627	0.547	0.299	0.102
Scenario C	167.2	0.461	0.606	0.458	0.302	0.108
July 1 start date, medium *P(T|C)*						

[26] In fact, the attack rate slightly increases in low-income and lower-middle-income regions because of interaction with seasonality effects: antiviral treatment in wealthier countries delays the pandemic slightly, which means that Northern Hemisphere countries (including India and China) have slightly higher average seasonality factors.

[27] The latter effect is primarily due to a reduction in attack rates in Russia, which constitutes a sizeable proportion of the upper-middle-income population and is at high latitude. As a result, it benefits from delaying the pandemic peak past winter.

Baseline (no AV)	0.0	0.547	0.546	0.536	0.617	0.538
Scenario A	117.3	0.456	0.544	0.532	0.268	0.110
Scenario B	138.8	0.358	0.542	0.331	0.249	0.102
Scenario C	167.2	0.301	0.484	0.183	0.183	0.077

Note: Results are for an AV efficacy of 0.6.

The small external benefits to high-income countries of countries in the lowest two income brackets having stockpiles still occur when the flu starts on July 1, assuming it is moderately transmissible. A highly transmissible flu, however, sharply reduces these benefits because a small stockpile is rapidly exhausted under a virus with a high reproductive rate. Indeed, in the January 1, high transmissibility case, high-income countries are slightly worse off in scenario C than scenario B. Lower-middle-income countries themselves see an additional 118 million additional cases relative to no treatment in this scenario because of a higher attack rate in China, where the slight delay from a 1 percent stockpile pushes the pandemic peak toward winter.

It seems unlikely that countries in the lowest two income brackets will purchase and maintain their own stockpiles of antivirus for use in the course of a pandemic. But the external benefits to high-income countries suggest the scope for a Pareto improvement: high-income countries may find it in their interest to fund the acquisition and distribution of antivirals in lower-income countries. As discussed more fully below, Table 7 implies that each antiviral dose administered in low-income or lower-middle-income countries reduces the number of influenza cases in high-income countries by one-third of a case or more when the transmissibility of the flu is low to moderate.

An alternative is for high-income countries to donate antivirals directly to the outbreak source. Consider a scenario in which high-income countries donate one-tenth of their stockpile directly to Indonesia.[28] Table 8 shows the effect of this policy when combined with a 0/1/5/10 scenario. In nearly every case, the donation policy reduces the global attack rate and the attack rate in high-income countries. In a few cases, attack rates are increased due to the seasonality effect. With the January 1 start date, slowing the pandemic spread increases the average seasonality factor in the Northern Hemisphere. The gain from donation is largest in the low

[28]Almost none of the high-income regions consume their entire 10 percent stockpile (assuming Moderate $P(T|C)$) at all or within one year, so the net effect of this policy is almost identical to one where high-income countries simply pay for the additional doses in Indonesia.

P(T|C) scenarios because in these cases, donation to the outbreak source reduces the R_0 below 1 and the pandemic is contained. The gains also are large with a July 1 start date, where delay is more valuable due to the interaction with seasonality. The external benefits of providing antivirals are lower when the antiviral efficacy is lower.

Table 8. Impact of Antivirals (AV) on One-Year Attack Rate, Scenario D

P(T|C)	Start date	AV efficacy	Income group	Baseline scenario	Scenario D	Difference
Low	January 1	0.6	World	0.178	0.001	0.176
Low	January 1	0.6	High	0.015	0.001	0.014
Medium	January 1	0.6	World	0.310	0.289	0.021
Medium	January 1	0.6	High	0.052	0.047	0.005
Medium	January 1	0.5	World	0.310	0.319	-0.009
Medium	January 1	0.5	High	0.052	0.068	-0.016
Medium	July 1	0.6	World	0.450	0.328	0.122
Medium	July 1	0.6	High	0.111	0.081	0.031
High	January 1	0.6	World	0.489	0.468	0.021
High	January 1	0.6	High	0.103	0.108	-0.005

Which strategy is more effective for high-income countries: paying for antiviral doses to be divided among lower-income countries in proportion to population, or paying for antiviral doses to be targeted to the outbreak country? Under what circumstances will welfare in high-income countries be increased by providing antivirals to lower-income countries? Table 9 compares the effectiveness of purchasing doses for lower-income countries in general to the effectiveness of purchasing doses for the outbreak source country. Neither strategy is effective in reducing the number of cases in high-income countries for a highly infectious pandemic (*P(T|C)* = 0.06).But in all other cases, targeting the outbreak source is dramatically more cost-effective than spreading doses throughout low-income or lower-middle-income regions.

Table 9. Number of High-Income Country Cases Reduced per Dose Purchased for Lower-Income Countries

Scenario transition	*P(T|C)*	Start date	Cases reduced
Scenario A ->Scenario B	Low	January 1	0.74
Scenario A ->Scenario B	Medium	January 1	0.31
Scenario A ->Scenario B	Medium	July 1	0.33
Scenario A ->Scenario B	High	January 1	0.00
Scenario B ->Scenario C	Low	January 1	0.29
Scenario B ->Scenario C	Medium	January 1	0.22
Scenario B ->Scenario C	Medium	July 1	0.83
Scenario B ->Scenario C	High	January 1	-0.17
Scenario B ->Scenario D	Low	January 1	1.42
Scenario B ->Scenario D	Medium	January 1	0.52
Scenario B ->Scenario D	Medium	July 1	3.06
Scenario B ->Scenario D	High	January 1	-0.51

Table 9 suggests that providing antivirals to developing countries in the event of a pandemic may pass a benefit–cost test. Although the percentage reduction in cases from providing antivirals is small, millions of cases of the flu in high-income countries would thereby be avoided at a cost of 3–4 doses of antivirals per case. Even at a cost of $25–$30 per course of treatment, this would likely pass a benefit–cost test, even without any fatalities: Sander et al. (2008) predict the economic cost of an influenza epidemic in the United States at $187 per person, based on a 50 percent attack rate. Keogh-Brown et al. (2010) suggest that a mild influenza pandemic in the United Kingdom similar to those in 1957 or 1968 would reduce GDP by 0.58 percent over the course of year; a more severe pandemic with a case fatality rate of 1 percent would reduce GDP by 4.5 percent over the course of a year. Typical estimates of the value of a statistical life in high-income countries are in the millions of dollars, so even a very low but positive fatality rate would lead to a large value from reducing cases.

Sensitivity of Results to Proportion Treated in Low-Income Countries

The effectiveness of antivirals in mitigating pandemics is a function of a country's ability to rapidly, accurately identify infectious individuals and provide them with treatment. In the primary simulation runs, we assume that any country that has a stockpile can treat 50 percent of symptomatic cases within a day. This assumption is arguably too optimistic, particularly for developing countries, even though it is more pessimistic than the assumptions Colizza et al.

(2007) propose. By weakening these assumptions and reducing the proportion of symptomatic infectious cases who receive treatment, we examine sensitivity to this assumption.

Table 10 presents simulations based on a scenario in which only 30 percent of symptomatic infectious cases can be treated in low-income and lower-middle-income countries, 40 percent in upper-middle-income countries, and 50 percent in high-income countries. The effect of reducing the proportion who are treated is dramatic, not only on the direct value of antiviral treatment in countries that have a stockpile, but also on the external benefits to high-income countries from providing antiviral doses to lower-income countries. In terms of direct effects, the marginal benefit (in terms of world cases reduced) from an antiviral stockpile in upper-middle- and high-income countries is reduced from 5.2 percentage points to 4.8 percentage points. The marginal benefit from adding a 1 percent stockpile to lower-middle-income countries is reduced from 8.4 percentage points to 2.5 percentage points. The marginal benefit from adding a 1 percent stockpile to low-income countries is reduced from 5.6 percentage points to 1.3 percentage points.

Table 10.Attack Rate after One Year, Standard versus Reduced Proportion of Infectious Treated

Income group		Strong health infrastructure	Weak health infrastructure
World	Baseline (no AV)	0.446	0.446
	Scenario A	0.394	0.398
	Scenario B	0.310	0.373
	Scenario C	0.254	0.360
Low	Baseline (no AV)	0.528	0.528
	Scenario A	0.529	0.529
	Scenario B	0.525	0.525
	Scenario C	0.458	0.518
Lower middle	Baseline (no AV)	0.401	0.401
	Scenario A	0.403	0.403
	Scenario B	0.169	0.336
	Scenario C	0.121	0.314
Upper middle	Baseline (no AV)	0.397	0.397
	Scenario A	0.210	0.254
	Scenario B	0.166	0.248
	1/1/5/10	0.051	0.224
High	Baseline (no AV)	0.325	0.325

	Scenario A	0.059	0.061
	Scenario B	0.052	0.058
	Scenario C	0.045	0.055

Note: Strong health infrastructure means 50 percent of new symptomatic infectious cases are treated in regions that have an antiviral (AV) stockpile; weak health infrastructure means that 30 percent of cases are treated in low-income and lower-middle-income regions, 40 percent in upper-middle-income regions, and 50 percent in high-income regions. All scenarios have a January 1 start date.

The benefits to high-income countries from treatment in low-income and lower-middle-income countries also are reduced significantly in the weak (compared to the strong) health infrastructure scenarios. The marginal reduction in the attack rate in high-income countries from funding a 1 percent stockpile for lower-middle-income countries falls from 0.7 percentage points to 0.3 percentage points. Similarly weak health infrastructure causes the benefits from funding a 1 percent stockpile to low-income countries to fall from 0.7 to 0.3. These results follow from a basic property of the core SIR model: with a lower percentage of people treated, the effect of antivirals on reducing the reproductive rate is diminished proportionally, which has a greater than proportional impact on the benefits of antivirals from reducing attack rates.

These results suggest that if the parameters of a weak health infrastructure more accurately describe the real world, the benefits to high-income countries of providing antivirals to lower-income countries are greatly reduced, a result that is confirmed by Table 11. But they also suggest that investments to increase the number of infectious people who can be treated in lower-income countries will have large benefits and will complement policies that provide antiviral doses to these countries.

Table 11. Number of High-income Country Cases Reduced per Antiviral Dose Purchased in Low-Income Countries, Sensitivity to Weak Health Infrastructure

Scenario transition	*P(T|C)*	Cases reduced	
		Strong health infrastructure	Weak health infrastructure
Scenario A ->B	Medium	0.31	0.12
Scenario B ->C	Medium	0.22	0.09
Scenario B ->D	Medium	0.52	0.21

Note: All scenarios have a January 1 start date.

7. Conclusions

Treating infectious individuals with antiviral drugs may be an effective method for mitigating the consequences of an influenza pandemic. While most of the benefits from a country choosing to treat its population with antivirals will accrue to that country, there are some positive externalities. This implies that in a world in which low-income countries are unlikely to hold stockpiles of antiviral drugs, it may be in the self-interest of wealthy countries to (collectively) fund the purchase and distribution of antivirals to low-income countries, even without any altruistic or humanitarian motivations.

The payoff to providing antiviral doses to low-income countries is illustrated by a simple, two-country SIR model. That model shows significant externalities and complementarities in antiviral treatment: when one country treats more of its population, it reduces the attack rate in the other country and increases the marginal benefit from additional treatment in the other country.

However, it would be misleading to draw policy conclusions from a simple two-country model. In reality, influenza spreads through a complex network of air- and land-based travel. The spread of the flu and the effectiveness of treatment policies depend on seasonal factors—e.g., whether the flu peaks in Northern Hemisphere winter or summer. In addition, the number and distribution of low-income versus high-income countries differs significantly from the symmetric two-region SIR model. We simulate the spread of the flu in a more descriptively realistic global epidemiological model in order to capture the impact of these features on the effectiveness of treatment policies.

Under our base case assumptions of moderate transmissibility of the flu, the distribution of antiviral stockpiles from high-income countries to low-income and lower-middle-income countries may indeed pay for itself: providing a stockpile equal to 1 percent of the population of these countries will reduce cases in high-income countries after a year by about 6.13 million cases at a cost of 4.62 doses per case avoided. Concentrating doses on the outbreak country is even more cost-effective: in our base case, this policy reduces the number of influenza cases in high-income countries by 4.76 million cases, at the cost of roughly 1.92 doses per case avoided.

These results are dependent on the transmissibility of the flu, its start date, the effectiveness of treatment in reducing infection, and the proportion of infectious people who can realistically be identified and treated. Our simulations reveal that reducing the proportion of symptomatic infectious that can be treated from our base case of 50 percent to 40 percent in lower-middle-income countries and 30 percent in low-income countries more than doubles the

number of doses required to reduce a case of the flu in high-income countries. Providing stockpiles to low-income countries may still pass the benefit–cost test, but our results suggest that improving the delivery of health services in low-income countries will complement policies to treat pandemic flu, in addition to yielding other health benefits.

References

Baroyan, O.V., G.A. Mironov, and L.A. Rvachev.1981. An Algorithm Modeling Global Epidemics of Mutant Origin. *Programming and Computer Software* 6(5):272–77.

Colizza, V., et al. 2007.Modeling the Worldwide Spread of Pandemic Influenza: Baseline Case and Containment Interventions. *PLoS Medicine* 4(1): e13.

Cooper, B.S., et al.2006.Delaying the International Spread of Pandemic Influenza. *PLoS Medicine* 3(6): e212

Epstein, J.M., et al.2007. Controlling Pandemic Flu: The Value of International Air Travel Restrictions. *PLoS ONE* 2(5): e401.

Ferguson, N.M., et al. 2005. Strategies for Containing an Emerging Influenza Pandemic in Southeast Asia. *Nature* 437(7056): 209–14.

Ferguson, N.M., et al.2006. Strategies for Mitigating an Influenza Pandemic. *Nature*, advance online publication, April 26, 2006.doi: 10.1038/nature04795.

Germann, T.C., et al.2006.Mitigation Strategies for Pandemic Influenza in the United States. *Proceedings of the National Academy of Sciences* 103(15): 5935–40.

Glezen, W.P. 1996. Emerging Infections: Pandemic Influenza. *Epidemiology Review* 18(1): 64–76.

Grais, R.F., J.H. Ellis, and G.E. Glass.2003. Assessing the Impact of Airline Travel on the Geographic Spread of Pandemic Influenza. *European Journal of Epidemiology* 18(11): 1065–72.

Guimerà, R., et al. 2005. The Worldwide Air Transportation Network: Anomalous Centrality, Community Structure, and Cities' Global Roles. *Proceedings of the National Academy of Sciences* 102(22): 7794–99.

Hajdin, C., et al. 2009. Stochastic Equation-Based Model of a Global Epidemic (Version 3.0).Research Triangle Park, NC: RTI International.

Hethcote, H.W., and J.A. Yorke. 1984. Gonorrhea Transmission Dynamics and Control. Lecture Notes in Biomathematics 56.Berlin, Germany: Springer.

Hufnagel, L., D. Brockmann, and T. Geisel.2004.Forecast and Control of Epidemics in a Globalized World. *Proceedings of the National Academy of Sciences* 101(42): 15124–29.

Keogh-Brown, M.R., et al. 2010.The Possible Macroeconomic Impact on the UK of an Influenza Pandemic. *Health Economics* 19(11):1345–60.

Kermack, W.O., and A.G. McKendrick.1927. A Contribution to the Mathematical Theory of Epidemics. *Proceedings of the Royal Society of London A* 115(772): 700–21.

Longini, I.M., Jr., et al.2005.Containing Pandemic Influenza at the Source. *Science* 309 (5737): 1083–87.

Lowen, A.C., et al. 2008. High Temperature (30°C) Blocks Aerosol but Not Contact Transmission of Influenza Virus. *Journal of Virology* 82(11): 5650–52.

Mossong, J., et al. 2008. Social Contacts and Mixing Patterns Relevant to the Spread of Infectious Diseases. *PLOS Medicine* 5(3): 381–91.

Over, M., and P. Piot. 1993. HIV Infection and Sexually Transmitted Disease. In *Disease Control Priorities in Developing Countries*, edited by D.T. Jamison et al. New York: Oxford University Press.

Rvachev, L.A., and I.M. Longini, Jr.1985.A Mathematical Model for the Global Spread of Influenza. *Mathematical Biosciences* 75(1): 3–22.

Sander, B., et al.2008. Economic Evaluation of Influenza Pandemic Mitigation

Strategies in the United States Using a Stochastic Microsimulation Transmission Model. *Value in Health* 12(2): 226–33.

Shaman, J. and Kohn, M. (2009). "Absolute humidity modulates influenza survival, transmission, and seasonality", *Proceedings of the National Academy of Sciences,* 106, pp. 3243-3248

United Nations.2007.*World Urbanization Prospects: The 2007 Revision*. New York.

Vynnycky, E., et al. 2007.Estimates of the Reproduction Numbers of Spanish Influenza Using Morbidity Data. *International Journal of Epidemiology* 36(4):881–89.

Vynnycky, E., and W.J. Edmunds.2008.Analyses of the 1957 (Asian) Influenza Pandemic in the United Kingdom and the Impact of School Closures. *Epidemiology and Infection*, 136: 166–79.

Appendix A. Results for Two-City Model, Sensitivity to Key Parameters

β^*	e	α	p_A	p_B	Attack rate in City A	Attack rate in City B
0.3	0.4	0.01	0	0	0.588	0.579
			0.6	0	0.301	0.563
			0	0.6	0.565	0.296
			0.6	0.6	0.240	0.235
0.35	0.4	0.01	0	0	0.719	0.707
			0.6	0	0.475	0.705
			0	0.6	0.709	0.468
			0.6	0.6	0.456	0.449
0.25	0.4	0.01	0	0	0.376	0.370
			0.6	0	0.111	0.326
			0	0.6	0.329	0.107
			0.6	0.6	0.012	0.005
0.3	0.3	0.01	0	0	0.588	0.579
			0.6	0	0.381	0.569
			0	0.6	0.573	0.376
			0.6	0.6	0.353	0.348
0.3	0.5	0.01	0	0	0.588	0.579
			0.6	0	0.224	0.555
			0	0.6	0.557	0.220
			0.6	0.6	0.090	0.083
0.3	0.4	0.01	0	0	0.588	0.579
			0.7	0	0.249	0.558
			0	0.7	0.559	0.244
			0.7	0.7	0.145	0.138
0.3	0.4	0.01	0	0	0.588	0.579
			0.5	0	0.355	0.567
			0	0.5	0.570	0.349
			0.5	0.5	0.318	0.313
0.3	0.5	0.015	0	0	0.587	0.580
			0.6	0	0.318	0.555
			0	0.6	0.557	0.314
			0.6	0.6	0.239	0.235
0.25	0.5	0.005	0	0	0.588	0.579
			0.6	0	0.278	0.572
			0	0.6	0.574	0.271
			0.6	0.6	0.240	0.234

$\beta°$=untreated rate of infection; e=the proportionate reduction in $\beta°$ achieved by the antiviral;α =fraction who travel; p_A= proportion of new infectious people receiving antiviral treatment in City A; p_B =proportion of new infectious people receiving antiviral treatment in City B.

Appendix B. Derivation of the Contact Rate Matrix

This section describes the derivation of the matrix of contact rates C_i based on data from Europe, which is adjusted to fit other cities by accounting for differences in age structure and population density.

The four-by-four contact rate matrix with elements C_{ijk} describes the number of contacts per day that an infectious person of age group j has with individuals in age group k in city i. Following Hethcote and Yorke (1984) and Over and Piot (1993), we decompose this matrix into two parts: a four-by-four mixing matrix M_i and a four-by-one contact vector A_i, such that

$$C_i = M_i \cdot I_4 \cdot A_i,$$

where I_4 is the four-by-four identity matrix.

The vector A_i gives the absolute number of contacts per day for each age group; for example, $A_{Jakarta}$'= [7.29, 9.45, 6.41, 3.98] means that in Jakarta, each infectious case with age group a=0 has 7.29 contacts per day, each infectious case with a=1 has 9.45 contacts per day, and so forth. We define the four-by-one vector n_i as containing the proportion of the population in each age category, so the elements of n_i sum to 1. The total number of contacts per person per day in city i is given by the product $n_i'A_i$. For example, in Indonesia, 9.64 percent of the population are in age group 0, 18.75 percent are in age group 1, 66.09 percent are in age group 2, and 5.52 percent are in age group 3, so the total contacts per day is the product $n_{Jakarta}'A_{Jakarta}$ = 6.94.

The mixing matrix M_i describes the relative mixing rates between age groups. Each row of the M matrix sums to 1, so the element M_{ijk} describes the proportion of an individual in age group j's contacts that occur with age group k. For example, if the j th row of M_i were [0.25, 0.25, 0.25, 0.25] then age group j individuals in city i would contact equal numbers of all four age groups each day.

C, M, and A have i subscripts because these vary across cities, although all cities within a region are assumed to have the same values of C, M, A and n. Values for A and M are generated based on mixing data from Mossong et al. (2008) and data on population density and

urbanization rates. Mossong et al. measure physical contact rate matrices for eight European countries and ten age categories from self-reported contacts across 7,290 individuals.[29] Their data provides the equivalent of a *C* matrix for each of the eight countries. The goal is to use these data to extract an underlying set of core behaviors that can be assumed to hold in all countries and will allow reconstruction of a *C* matrix for every region in the Global Epidemiological Model.

We use the Mossong data to estimate the absolute number of contacts in each city and the relative mixing rates of one age group with another. *A* vectors are defined for each region such that $n_i'A_i = 6.94$ in every region, where 6.94 is the weighted-average absolute contact per day across the eight Mossong countries and across all age groups.[30]The *M* matrix reflects the age structure of each area; hence it is inappropriate to use the *M* matrices reported in Mossong directly in the Global Epidemiological Model.

We estimate an *M* matrix for each region of the model by assuming a structural form for the matrix and then using the eight Mossong matrices to estimate an underlying mixing matrix stripped of the country-specific *n* vectors. Following Hethcote and Yorke (1984), we define a four-by-one vector B_i by

$$B_{ij} = \frac{A_{ij}n_{ij}}{\sum_{m=1}^{4} A_{im}n_{im}},$$

and a four-by-four matrix G by

$$B_{ij}\left(1 - G_{ijk}\right) = M_{ijk},$$

where G_{ijk} is the *(j,k)* element of the *B* matrix for city *i*, and M_{ijk} is the *(j,k)* element of the *M* matrix for city *i*. B_{ij} is the *j* th element of the *B* vector in city *i* and can be interpreted as the proportion of contacts that would be with age group *j* city *i* under random mixing. The *G* matrix allows us to move away from random mixing and captures the more complex structure where people can have differential mixing rates across age groups.

---

[29] The countries are Belgium, Finland, Germany, Great Britain, Italy, Luxembourg, the Netherlands, and Poland.

[30] This crucial assumption, made for simplicity, implies that total contacts per day do not vary by age structure; it will not be the case that the pandemic spreads faster in countries with younger populations merely because young people have higher contact rates. An alternative structure that would have this property could be generated by fixing the individual elements A_i across countries, but that would lead to very large variation in contact rates (and thus attack rates) across countries.

We assume a structural form for the *G* matrix where $G_{ijk} = G_{ij}$, i.e., all elements of the same row have the same value.[31] In other words, we assume each age group has the same relative preference for mixing with its own age group as for all other age groups. We estimate the *G* matrix that minimizes the sum of squared residuals across the eight Mossong countries to find an ―optimal" *G* matrix, *G**. Then we reconstruct *M* and *C* matrices for every region in the model by substituting *G** and the region-specific age structures from real-world population data into the equations above.

We adjust contact rates for population density, recognizing that areas with higher population density have higher contact rates. We assign every city and rural area to one of four categories: relative density d=1 for rural areas with residual urbanization less than 40 percent, d=1.1 for rural areas with residual urbanization greater than 40 percent; d = 1.3 for cities with density less than 1,000 people per square kilometer; and d=1.4 for cities with density greater than 1,000 people per square kilometer. In general, rural areas in high-income countries have high residual urbanization, while those in low-income countries have low residual urbanization. Cities in high-income countries have low population densities, while those in low-income countries have high population densities.

Appendix C. List of Regions in the Global Epidemiological Model, by Income Group

Regions in the Low-Income Group:

Bangladesh, Bolivia, Cambodia, Egypt, Georgia, India, Indonesia, Kyrgyz Republic, Madagascar, Malawi, Morocco, Mozambique, Nicaragua, Nigeria, Pakistan, Paraguay, Philippines, Senegal, Sri Lanka, Tanzania, Uganda, Viet Nam, Zambia, Zimbabwe, Rest of Central Africa, Rest of East Asia, Rest of Eastern Africa, Rest of Former Soviet Union, Rest of Oceania, Rest of South African Customs Union, Rest of South America, Rest of South Asia, Rest of South-Central Africa, Rest of Southeast Asia, Rest of Western Africa

[31] Tested alternative forms showed that moving from a 16-element *G* matrix to a four-element one had very little impact on the fit to the Mossong data but moving from a four-element G matrix to a one-element G matrix significantly worsened the fit.

Regions in the Lower-Middle-Income Group:

Albania, Argentina, Armenia, Azerbaijan, Botswana, Brazil, Bulgaria, China, Colombia, Ecuador, Iran, Kazakhstan, Malaysia, Mauritius, Peru, Romania, South Africa, Thailand, Tunisia, Ukraine, Rest of Central America, Rest of Eastern Europe, Rest of Europe, Rest of North Africa, Rest of Western Asia

Regions in the Upper-Middle-Income Group:

Chile, Croatia, Estonia, Hungary, Latvia, Lithuania, Mexico, Poland, Republic of Korea, Russian Federation, Slovak Republic, Taiwan, Turkey, Uruguay, Venezuela, Rest of the Caribbean

Regions in the High-income Income Group:

Australia, Austria, Belgium, Canada, Cyprus, Czech Republic, Denmark, Finland, France, Germany, Greece, Hong Kong Special Administrative Region of China, Iceland, Ireland, Israel and Arabia, Italy, Japan, Luxembourg, Malta, Netherlands, New Zealand, Portugal, Singapore, Slovenia, Spain, Sweden, Switzerland, United Kingdom, United States of America, Rest of European Free Trade Association, Rest of North America