

NBER WORKING PAPER SERIES

POLICY RESPONSE TO PANDEMIC INFLUENZA:
THE VALUE OF COLLECTIVE ACTION

Georgiy Bobashev
Maureen L. Cropper
Joshua M. Epstein
D. Michael Goedecke
Stephen Hutton
Mead Over

Working Paper 17195
<http://www.nber.org/papers/w17195>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2011

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. The findings and conclusions of this paper are those of the authors and do not necessarily represent the views of the National Bureau of Economic Research, the World Bank and its affiliated organizations, the Executive Directors of the World Bank, or the governments they represent.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Georgiy Bobashev, Maureen L. Cropper, Joshua M. Epstein, D. Michael Goedecke, Stephen Hutton, and Mead Over. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Policy Response to Pandemic Influenza: The Value of Collective Action
Georgiy Bobashev, Maureen L. Cropper, Joshua M. Epstein, D. Michael Goedecke, Stephen Hutton, and Mead Over
NBER Working Paper No. 17195
July 2011
JEL No. D62,I18

ABSTRACT

This paper examines positive externalities and complementarities in the use of antiviral pharmaceuticals to mitigate pandemic influenza. The paper demonstrates the presence of treatment externalities in simple epidemiological SIR models, and then through simulations of a Global Epidemiological Model, in which the pandemic spreads between cities through the international airline network, and between cities and rural areas through ground transport. While most treatment benefits are private, spillovers may mean that it is in the self-interest of rich countries to pay for some AV treatment in poor countries. The most cost-effective policy is for rich countries to donate doses to the outbreak source country; however, donating doses to poor countries in proportion to their populations may also be cost-effective. These results depend on the transmissibility of the flu strain, the efficacy of antivirals in reducing transmissibility and on the proportion of infectious that can be identified and treated.

Georgiy Bobashev
RTI International
3040 E. Cornwallis Road
Research Triangle Park, NC 27709-2194
bobashev@rti.org

D. Michael Goedecke
RTI International
3040 E. Cornwallis Road
Research Triangle Park, NC 27709-2194
mgoedecke@rti.org

Maureen L. Cropper
Department of Economics
University of Maryland
College Park, MD 20742
and NBER
cropper@econ.umd.edu

Stephen Hutton
The World Bank
1818 H Street NW, MSC 9-001
Washington, DC 20433
shutton@worldbank.org

Joshua M. Epstein
Johns Hopkins University School of Medicine
733 North Broadway
Baltimore, MD 21205
jepste15@jhmi.edu

Mead Over
Center for Global Development
mover@cgdev.org

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 which had significant economic costs, as well as causing millions of deaths and illnesses.¹ The 1918 pandemic killed an estimated 3% of the world's population, despite being slow to spread in an age before 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 the continued risks to the world should a pandemic occur. They also remind us that there is relatively little cooperation and coordination between countries, and that wealthy nations prioritize 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 to slow the spread of the flu in one county reduce attack rates in other countries? Does treatment of infected persons in one country increase the marginal benefits of treatment policies in other countries (i.e., are there treatment complementarities?). 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: (1) Is it cost-effective for wealthy nations to pay for the purchase and distribution of antivirals in poor countries to slow the spread of the pandemic? (2) What global allocation rules are most effective in reducing attack rates? In the case where treatment with antiviral drugs alone can contain a potential pandemic, there is an obvious case for wealthy nations to pay for pandemic containment. 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 such a scheme: 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?

¹ These were the 1918, 1957 and 1968 pandemics.

Our Approach

To investigate these issues we use a detailed Global Epidemiological Model to simulate influenza pandemics under a range of conditions and antiviral treatment policies. The model divides the world into 106 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 a rural area in each region via land travel. Within each city or rural area the flu spreads via a model in which people transition from susceptible to exposed to infectious to recovered or dead. 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 (it decreases as one moves away from the equator) and season (it 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. Poor countries are assumed to have no stockpiles. We ask whether it pays rich countries to pay for stockpiles to be distributed in poor countries. We compare two rules for rich countries distributing stockpiles in poor 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 poor country. We judge the success of a control strategy in terms of its impact on the attack rate (percent of the population infected) at the end of a year. Does it (collectively) pay rich 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 (influenza that peaks in Northern Hemisphere winter infects more people) and on the infectiousness of the flu (i.e., the reproductive rate R_0 , which measures the number of people an infectious person would infect in an otherwise totally susceptible population). 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 a pandemic can be contained in the source country (Ferguson, et al., 2005; Longini, et al., 2005). In the more likely case in which the pandemic will spread through air travel the externalities associated with anti-viral treatment are smaller: 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.

It is always in the interest of wealthy countries to purchase and distribute antiviral doses in the outbreak country when doing so can contain a pandemic. In other cases, the marginal private benefits from using limited antiviral supplies to treat domestic patients exceed the marginal benefits from donating those doses abroad. But, if wealthy countries retain a stockpile of antiviral drugs sufficient to treat their own cases, then they can increase their welfare by paying for purchase and distribution of additional doses to the outbreak source country.

The benefits to rich countries of paying for antivirals in poor countries under "midrange" assumptions of influenza transmissibility cover the costs: donation of antiviral to the outbreak source country reduces the number of influenza cases in rich countries after 1 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 at a zero percent case fatality rate; 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 S-I-R 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 the 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 Global Epidemiological Model that 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 both models with significant detail at the country level, such as the models of influenza transmission in Thailand and Southeast Asia (Ferguson et al. 2005; Longini et al. 2005) and the U.S. and the U.K. (Ferguson et al. 2006; Germann et al. 2006) and global models of influenza transmission (Rvachev and Longini 1985; Grais, Ellis and Glass 2003; Colizza et al. 2007; Epstein et al. 2007). 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 anti-viral drugs—distributing drugs to all persons in the same school or workplace as an infected person, or, alternately, ring prophylaxis.

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. Models of global influenza transmission have focused on the role of the air transport network in the spread of influenza. 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 (S-E-I-R) model. The subsequent literature (Grais, Ellis and Glass 2003; Hufnagel et al. 2004; Cooper et al. 2006; Colizza et al. 2007, Epstein et al. 2007) has used a more detailed characterization of the air transport network (Colizza et al. (2007) model 3,500 airports) and has introduced randomness in travel rates and disease transmission. These global models have examined the impact of air travel restrictions and the effects of local policies to reduce disease transmission, but have not examined the economics of pandemic flu control, nor have they 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 Susceptible-Infectious-Recovered (S-I-R) model. Individuals are either Susceptible, i.e., capable of being infected, Infectious, i.e., capable of infecting others, 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:

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

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


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

β represents the probability of transmission conditional on exposure to an infected person, or the average rate of infection per susceptible. 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 5 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 susceptibles falls below $1/R_0$, and steadily declines thereafter. The epidemic is effectively over in less than 4 months, with only 0.01N new infectious cases

occurring after day 104, where N is the population size. The attack rate² $R(\infty)$, is approximately 58%.

Figure 1: Progression of Influenza in a One-City S-I-R Model



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, β° :

$$(4) \quad \beta = \beta^\circ(1-pe),$$

where e is the proportionate reduction in β° achieved by the anti-viral (e.g., $e=0.4$ implies a 40% reduction in the infectiousness of a sick person). Higher values of p reduce infectiousness and, hence, R_0 .

Assume that the government chooses the proportion of infectious 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 maximize³

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

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

³ This assumes that the decision maker focuses on the attack rate. It is also true that antivirals could reduce the duration of illness for those who become infected. This benefit is ignored here.

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

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

where $R_0 = \beta^0(1-pe)/\delta$.

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) call for p to be chosen to the point where the marginal benefits of increasing $S(\infty)$ both in terms of reducing the number never infected and the number of doses needed to treat the infectious, equal or exceed the cost of an additional dose,

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

The second-order conditions for an interior maximum, $d^2F(p)/dp^2 < 0$, are, however, 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 greater than 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 $\beta^0 = 0.3$, $\delta = 0.2$ 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 above example ($\beta^0 = 0.3$, $\delta = 0.2$ 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 is usually assumed that $p \leq 0.6$ (Longini et al. 2005; Germann et al. 2006). If $p=1$ is optimal in the unconstrained problem then $p = p^*$ and the number of antiviral doses that is used will be $p^*R(\infty)N$. In our example with $\beta^0 = 0.3$, $\delta = 0.2$ and $e = 0.2$, R_0 is reduced to 1.32 if $p=0.6$ (rather than 1.2 when $p=1$). The attack rate will be $R(\infty) = 0.45$ and a stockpile = 27% of the population will be used. If $e = 0.4$, $R(\infty) = 0.24$ and the

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 (Vynnycky and Edmunds 2008). Vynnycky et al. (2007) estimate an R_0 of 2.4-4.3 for the 1918 Spanish flu.

stockpile that will be used is approximately 14.4% of the population.

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 an S-I-R 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

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

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

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

$$(11) \quad \beta_A = \beta_A^0 (1 - p_A e)$$

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 travel between cities and that the same number of susceptibles also travel, in order to keep the population constant.

Figure 2: Spread of the Flu in a Two-City SIR Model

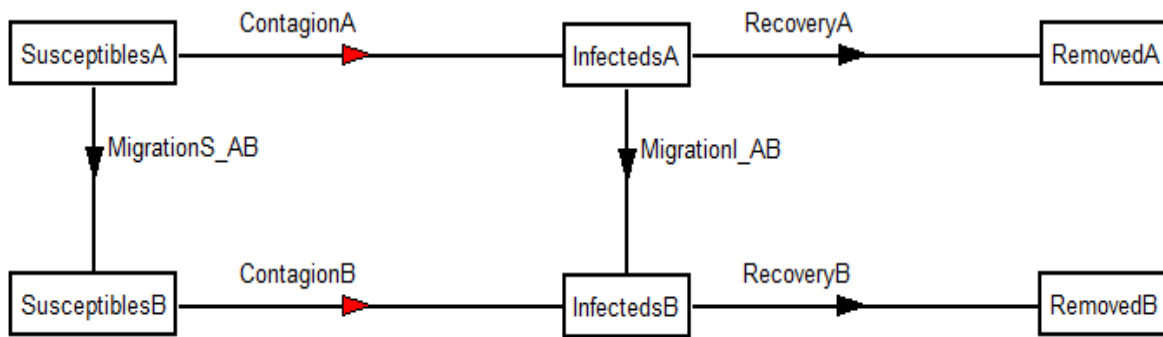



Figure 3 illustrates disease dynamics for the two cities assuming that the flu starts in city A ($I_A(0) = .001$) and moves to city B. Figure 3 assumes that neither city uses antivirals ($p_A = p_B = 0$) and that $\beta^0 = 0.3, \delta = 0.2$ and $\alpha = 0.01$. The number of infectious peaks first in the outbreak city (city A), but the attack rates are approximately equal (58%) in the two cities.

Figure 3: Progression of Influenza in a Two-City S-I-R 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% of its infectious or none.

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

β^0	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

Three points about Table 1 deserve emphasis: (1) 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; (2) when only one city treats its infectious, the external benefits to the no-treatment city are small;

(3) when both cities treat, the effectiveness of each city's stockpile is greater than when only one city employs a treatment strategy. When city A treats 60% of its infectious in isolation, the attack rate is reduced from 58% to 24%. In the two-city case, the attack rate is reduced to 30%. People from city B continue to re-infect city A, making city A's treatment of its infectious less effective. The benefits to city B of A's stockpile are, however, small: when city B has no stockpile, its attack rate is reduced by only 1.6 percentage points by virtue of the fact that A treats its infectious. When both cities treat their infectious, the benefits to both cities increase, compared to the case in which each city acts in isolation. When both cities treat, the attack rate is reduced from about 58% to 24% 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.)

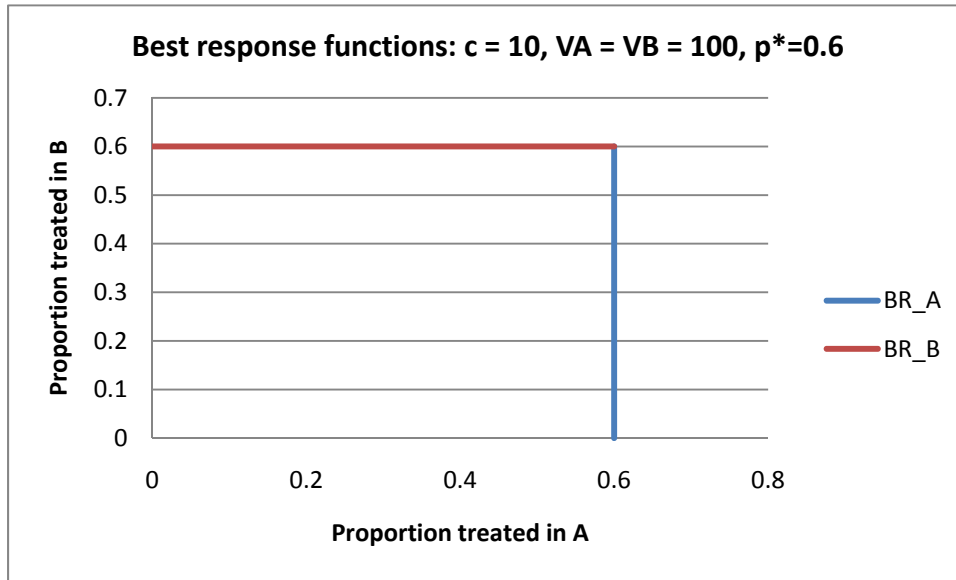
The magnitude of the externalities and complementarities observed in Table 1 are affected by the transmissibility of the flu (β^0/δ), the proportion of infectious who can be treated (p^*) and by the effectiveness of treatment (e). The benefits to the no-treatment city of antivirals in the treatment city are greater the less transmissible is the flu, as are the complementarities of treatment (see Table 1). Treatment complementarities are also greater the higher the proportion of infectious who can be treated (p^*) and the greater the efficacy of treatment (the higher is e). (See Appendix A.)

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 $\beta^0 = 0.3$, $\delta = 0.2$, $e = 0.4$, $\alpha = 0.01$ 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 "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.

Figure 4: Nash equilibrium, high value scenario




Suppose now that city A is a poor 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 poor city and rich city are pictured in Figure 5. The best response for city A depends on the value of p_B ; at low values of p_B city A will choose $p_A = 0$, but at high values of p_B city A will choose $p_A = 0.6$, and the Nash equilibrium remains at $p_A = p_B = 0.6$.

Figure 5: Nash equilibrium, City A decision depends on AV in City B



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 picture 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 $p_A = p_B = 0.6$. It is also the case that it pays city B to pay for antivirals sufficient to treat 60% of infectious 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 ($100 * 0.061N$).

Figure 6: Nash equilibrium, Low value in city A



Whether it pays city B to pay for antiviral treatment in city A depends, however, on the transmissibility of the flu, the efficacy of treatment and the maximum proportion of infectious that can be treated. The benefits to City B of paying for antiviral treatment in City A do not pass the benefit-cost test for a more readily transmissible flu (the $\beta^o = 0.35$ in Table 1) when $V_B = 100$ and $c = 10$, and the benefits are lower 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. In order to simulate realistic treatment policies we have expanded and modified the Global Epidemiological Model (GEM) developed by Bobashev and Goedecke (Hajdin et al. 2009). The GEM is a discrete time, stochastic S-E-I-R 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, Mironov, and Rvachev (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.⁸ Every region contains at least one city, and the network is fully connected.⁹ A rural area exists for every region except small islands and city states. Rural areas are defined as everything that is not explicitly urban, so the “rural” areas contain the population of all cities in the world that are not explicitly included in the model. Rural areas are treated the same as cities in the model, except that they are connected only to the cities within the same region, and not to cities in other regions. (Total world population is roughly 6.4 billion, of which roughly 84% is allocated to a rural area.) Cities are chosen to include the 130 most populous cities in the world and the 130 cities with the highest average airport flows per day.¹⁰

The cities and rural areas are connected by a daily flow between every city and its associated rural area, and by airline flows between cities. It is assumed that 1% of the population of each urban area travels to its associated rural area each day, and an equal number travel from the rural area to the urban area.¹¹ This assumption is somewhat arbitrary, but 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. Travelers are randomly selected 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 (once symptomatic infectious and dead cases are excluded).

Travel and population data come from a variety of sources. Data on airline flows are drawn from Official Airline Guide (OAG) statistics on flight schedules provided by L. Amaral (Guimerà, Mossa, Turtschi, and Amaral 2005). Numbers of passengers are based on the daily

⁷ Though the GEM is stochastic, the simulations reported here are run using a deterministic version of the model. There was little variation in the attack rate (our primary outcome variable) across runs of the stochastic version of the model.

⁸ Roughly 86% 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 This means that we end up with many cities in rich countries, and relatively fewer cities in poorer countries. Runs of the model excluding some rich country cities produced results similar to those reported below.

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

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

seat capacity between each pair of cities. The raw data track only direct flights, but a proportion of travelers is assumed to take 2-leg, indirect flights.¹³ Data on regional populations come from the World Bank development data portal. Data on urban metropolitan area populations are drawn primarily from the United Nations World Urbanization Prospects (2007), which uses broad definitions of urban areas.¹⁴ Rural area populations are calculated by deducting urban populations from regional populations.

Disease spread dynamics

Within each city and rural area, the progression of the flu is described by an S-E-I-R model.¹⁵ Individuals transition from being susceptible (S) to exposed (infected) (E) to infectious (I) (either symptomatic or asymptomatic) to recovered (R) or dead (D). Exposed individuals exhibit no symptoms (their infection is latent) and are not capable of infecting others, but will eventually become infectious (i.e., capable of transmitting the infection). Individuals who are infected and infectious are distinguished according to the number of days that have elapsed since they were first exposed (τ), and by whether they are symptomatic ($s=1$) or asymptomatic ($s=0$) (asymptomatic cases are less infectious). Thus $E(\tau, t)$ denotes the number of people who are Exposed on day t who became infected on day $t-\tau$. $I(s, \tau, t)$ denotes the number of people who are infectious on day t who first became infected on day $t-\tau$. The maximum length of the latent period is τ_1 days; the maximum length of time between first being infected and recovering or dying is τ_2 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.)

$$(12) \quad N(t) = S(t) + \sum_{\tau=0}^{\tau_1} E(\tau, t) + \sum_{\tau=1}^{\tau_2} \sum_{s=0}^1 I(s, \tau, t) + R(t) + D(t)$$

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 $\sigma(\tau)$ as the probability that an Infectious person recovers. We assume for simplicity (following Longini et al. 2005) that 80% of Exposed cases transition to Infectious after 1 day and 20% after 2 days; they will become symptomatic-Infectious with probability 0.67 and asymptomatic-Infectious with probability 0.33. Infectious cases either

¹³ A dataset of 10% of domestic US airline ticket coupons for 2004 is used 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, 68.2% of tickets are assumed to be 1-leg, and 31.8% are assumed to be 2-leg (though this varies across each city pair in the model).

¹⁴ Some cities are urban agglomerations of multiple cities, like Washington DC-Baltimore.

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

recover or die after 5-8 days since they became Exposed, with 30% recovering after 5 days, 40% after 6 days, 20% after 7 days, and 10% after 8 days. This means $\tau_1 = 1$ and $\tau_2 = 7$.

Thus:

$$(13) \quad \gamma(\tau) = \begin{cases} 0.8 & \text{if } \tau = 0 \\ 1 & \text{if } \tau = 1 \\ 0 & \text{otherwise} \end{cases}$$

$$(14) \quad \sigma(\tau) = \begin{cases} 0.3 & \text{if } \tau = 4 \\ \frac{0.4}{0.7} & \text{if } \tau = 5 \\ \frac{0.2}{0.3} & \text{if } \tau = 6 \\ 1 & \text{if } \tau = 7 \\ 0 & \text{otherwise} \end{cases}$$

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

$$(15) \quad I(1, \tau, t + 1) = I(1, \tau - 1, t)(1 - \sigma(\tau - 1)) + (0.67) \sum_{m=1}^{\tau_2} [E(m, t)\gamma(m)]$$

$$(16) \quad I(0, \tau, t + 1) = I(0, \tau - 1, t)(1 - \sigma(\tau - 1)) + (0.33) \sum_{m=1}^{\tau_2} [E(m, t)\gamma(m)]$$

for symptomatic and asymptomatic cases, respectively.

It is assumed 97.5% of symptomatic cases recover and 2.5% die, while 98.75% of asymptomatic cases recover and 1.25% die.¹⁶ This implies that the number of Recovered cases evolves according to:

$$(17) \quad R(t + 1) = R(t) + \sum_{\tau=1}^{\tau_2} \sum_{s=0}^1 [(0.975s + 0.9875(1 - s))I(s, \tau, t)\sigma(\tau)]$$

while the number of Dead cases evolves according to:

$$(18) \quad D(t + 1) = D(t) + \sum_{\tau=1}^{\tau_2} \sum_{s=0}^1 (0.025s + 0.0125(1 - s))I(s, \tau, t)\sigma(\tau)$$

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 new Exposed

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

cases caused by one Infectious person on day t is proportional to the number of Susceptible individuals in the city and to the daily infectious contact rate λ , which is equivalent to the parameter β 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:

$$(19) \quad E(\tau, t + 1) = \begin{cases} \lambda f(t) \frac{S(t)}{N(t) - D(t)} \sum_{\tau=1}^{\tau_2} \sum_{s=0}^1 [I(s, \tau, t) \cdot r(s)], & \text{for } \tau = 0 \\ E(\tau - 1, t)(1 - \gamma(\tau - 1)), & \text{for } \tau = 1, \dots, \tau_1 \end{cases}$$

The preceding describes an S-E-I-R model with no air travel. In the model, the evolution of $S(\cdot)$, $E(\cdot)$, $I(\cdot)$, $R(\cdot)$ and $D(\cdot)$ 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:

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

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

Influenza transmissibility

The GEM differs from the model described above by distinguishing four population age groups. Each city is divided into 4 age groups: 0 – 4 years (group 0), 5 – 14 years old (group 1), 15 – 64 years old (group 2), and ≥ 65 years old (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

¹⁷ This means that the expected number of travelers moving from i to j will be less than X_{ij} , as long as there are any cases who are symptomatic infectious or dead, since these people do not travel.

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.

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

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 children contact mostly other school children, adults contact mostly other adults, and so forth. 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, it is assumed 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. Contact rates are adjusted across cities by urban density (for cities) and by the urbanization rate of the rural areas (see Appendix B for details). This has the effect of increasing contact rates (and thus attack rates) for cities in lower income countries (which tend to have high density cities), and for rural areas in higher income countries (which have higher urbanization rates).

The probability of transmission given contact does not vary with age, and is chosen to determine the reproductive rate of the virus (determined by the particular biological characteristics of the specific influenza strain).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 $f_i(t)$ (see equation (19)) reflects the fact that influenza infection rates are much higher in winter months than in summer in temperate zones, but that there is less seasonal variation in tropical or arctic conditions. The mechanism for this is not well understood. Biological mechanisms include the possibility that the influenza virus is more fragile in warmer temperatures and has lower transmission at higher humidity (confirmed in controlled environment tests (e.g., Lowen et al. 2008), perhaps because virus particles are coated in moisture and less able to infect new hosts).¹⁹ All global influenza models attempt to account for this; most models use some kind of step function as an approximation (see, for example, Rvachev and Longini 1985; Grais et al. 2003, Colizza et al. 2007, Cooper et al. 2006).

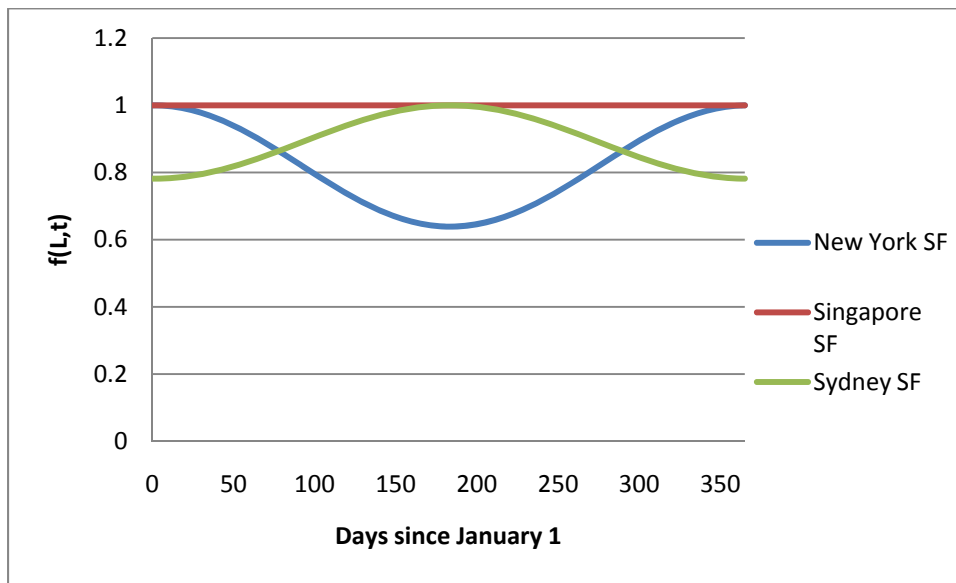
¹⁸ A simple model extension would be to allow $P(T|C)$ to vary by age-groups, so that for example a child or elderly individual might be more likely to contract the virus than a healthy adult. But reliable data on relative vulnerabilities is not available, and such a model change would be functionally 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.

The GEM incorporates seasonality through a seasonal infectiousness factor that varies by day of the year and by latitude of the city.²⁰ Infectiousness is assumed to be constant in the tropics, and elsewhere to follow 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).

Figure 7 shows the seasonality factors for Singapore (in the tropics), New York (temperate, Northern Hemisphere) and Sydney (temperate, Southern Hemisphere) with a pandemic starting on January 1 (winter in the Northern Hemisphere). Note that New York is further north than Sydney is south, and 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.

Figure 7: Seasonality factor for selected cities



Model Runs with No Treatment

In our base case the pandemic is assumed to begin with 100 infectious cases on January 1 in Jakarta. A pandemic influenza strain that mutates from an avian or other animal strain is more likely to occur in a country where people in rural areas live in close proximity to livestock. Jakarta is selected as a city in a large poor country that is well-connected to the international airline network. Although the model is run at a daily time-step for a 3-year period, we focus on the attack rate after 1 year. This represents the time that it might take to develop, test and mass produce a pandemic-strain-specific vaccine. Many of the gains from antivirus policy come from

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

delaying the spread of the virus, and delay is valuable in part because it buys time for vaccine development.

Every region in the model is assigned to an Income Group—Poor, Lower Middle, Upper Middle or Rich (see Table 2). Income group classifications are based on International Monetary Fund data on nominal 2008 GDP per capita, at 2008 exchange rates. Poor countries are those with GDP per capita less than \$3,000 USD, Lower Income are those with GDP per capita up to \$10,000, Upper Middle have incomes up to \$20,000 and Rich countries are those with per capita incomes exceeding \$20,000. For regions with multiple countries, the median income country is used to classify the region. We use this classification (rather than World Bank income group classifications) because we feel it is more representative of public health sector 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 group and age-group

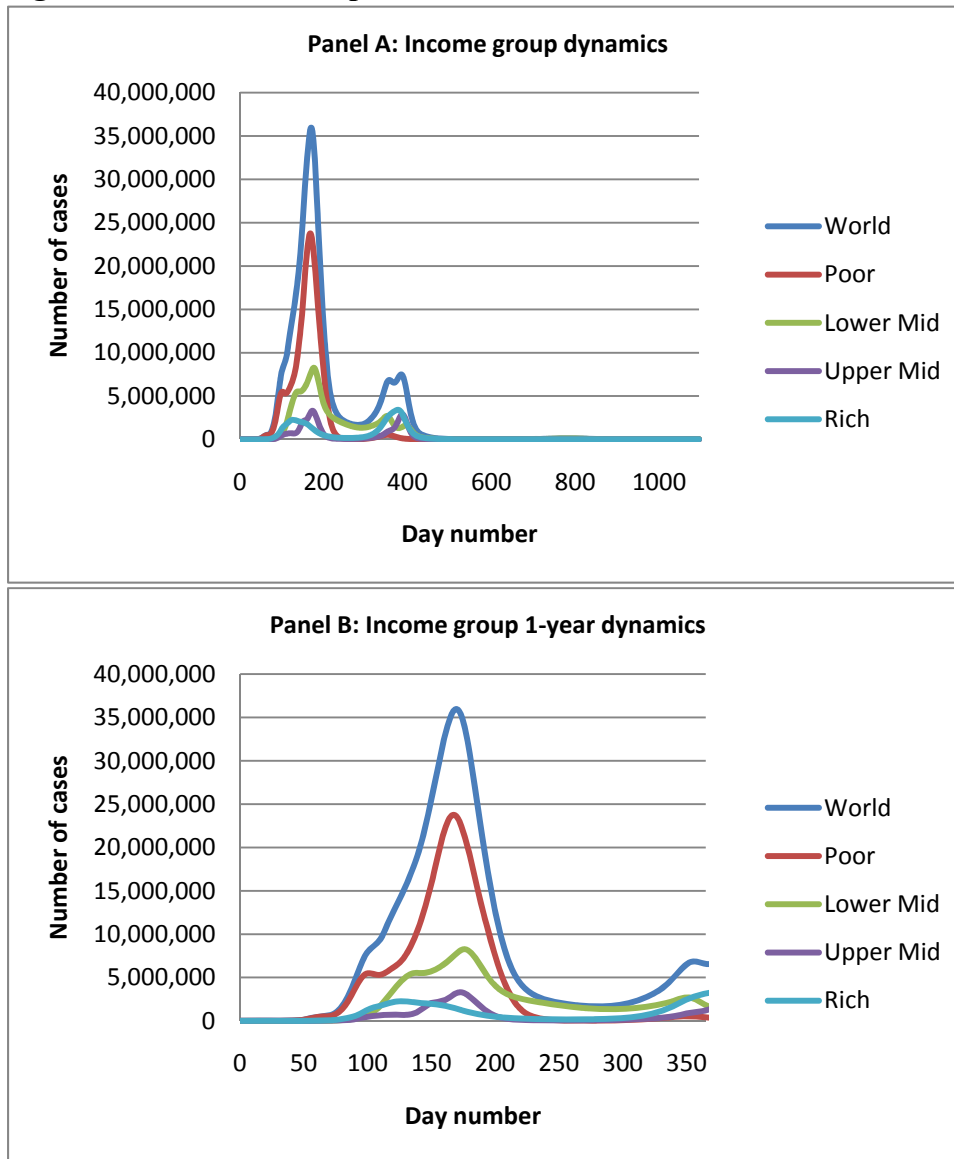
Income Group	Population				
	Total	Ages 0-4	Ages 5-14	Ages 15-64	Ages 65+
Poor	2,839,998,811	352,152,534	647,655,318	1,716,456,775	123,734,183
Lower Middle	2,156,161,724	165,680,600	361,106,698	1,473,910,542	155,463,884
Upper Middle	515,385,759	36,733,909	79,318,428	349,989,705	49,343,717
Rich	914,861,650	53,135,831	109,758,061	610,674,513	141,293,245
TOTALS	6,426,407,944	607,702,874	1,197,838,506	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% at the end of one year in the no-intervention case. In a global model the number of people that an infectious person infects in an otherwise susceptible population (R_0) varies from one country to another, making direct comparisons between $P(T|C)$ and R_0 in the models of sections 3 and 4 difficult.

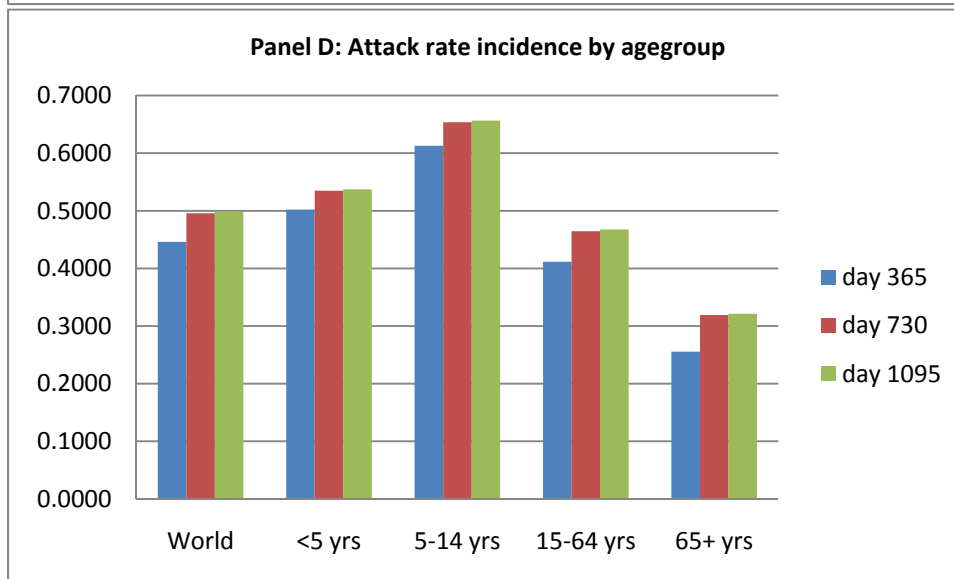

Figure 8 shows how the pandemic progresses in the no-treatment scenario. Panel A shows the number of new infectious cases on each day, by income group, for the duration of the pandemic, while Panel B shows the daily number of new infectious cases for the first year of the pandemic. Panel C shows attack 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 world-wide peak occurs around day 200, with a secondary wave of infection occurring in the second year, peaking around day 400. The second occurs as the Northern Hemisphere (where most of the world's population lives) re-enters winter. Most cases occur within a year; the 1-year attack rate is 44.6%, rising to 49.9% after 3 years.

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

Figure 8: Pandemic time path, no AV

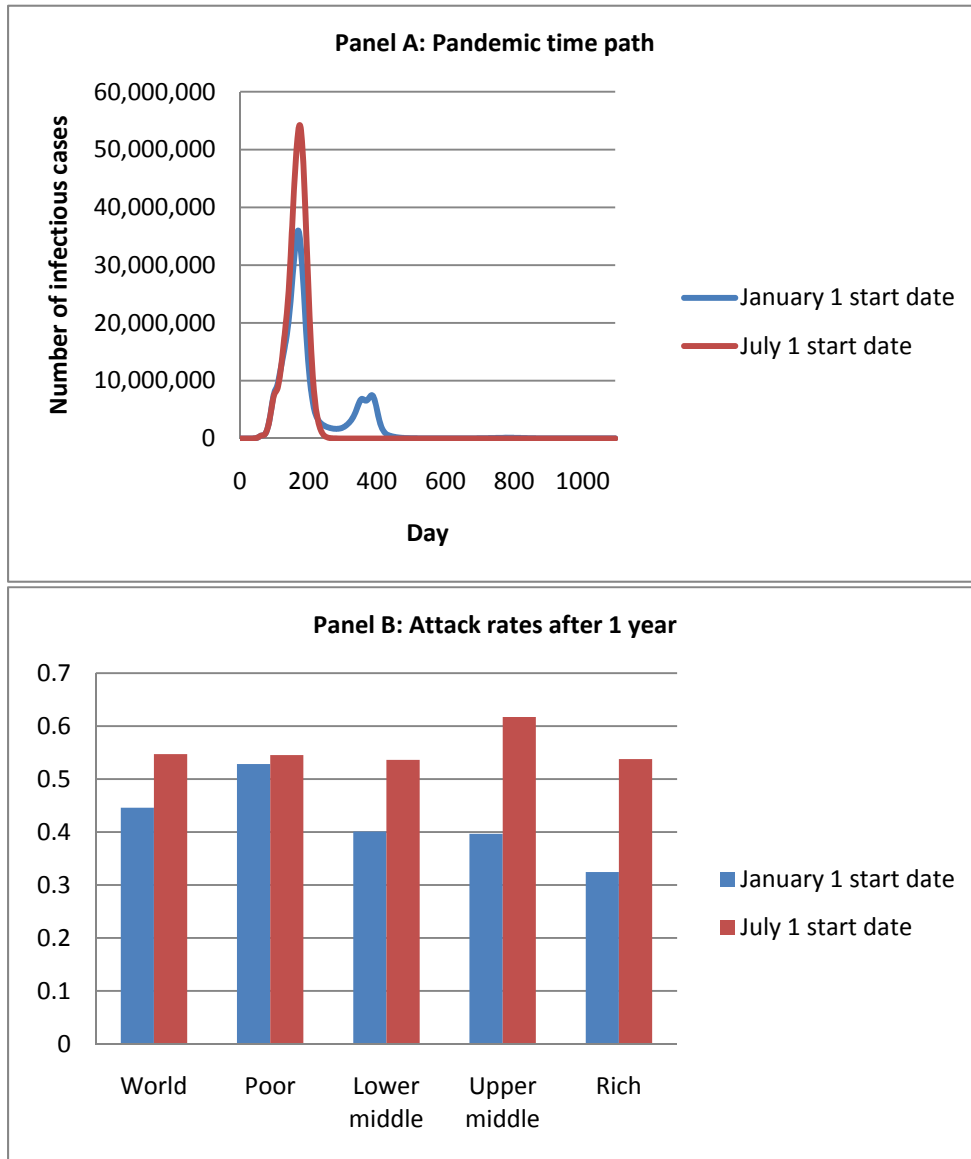


²¹ These patterns match data on attack rates in the United States during the 3 major pandemics of the 20th century (Glezen 1996).



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 (where the pandemic peaks during Northern Hemisphere summer) as compared to a July 1 start date (where the pandemic peaks during Northern Hemisphere winter). 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 and so the pandemic reaches a single, more intense peak. The increase in the 1-year attack rate is largest for rich countries, due to their location.

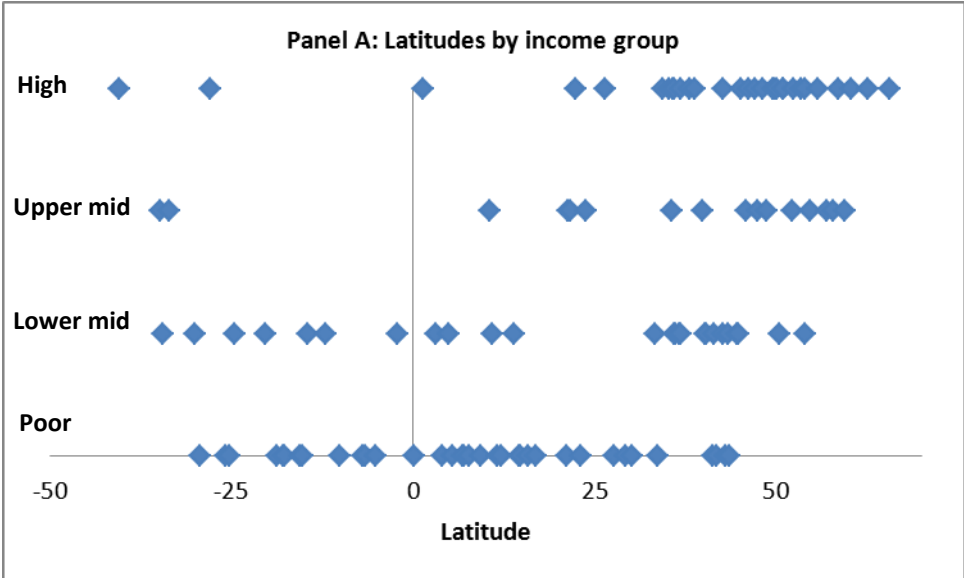
Figure 9: Pandemic dynamics, sensitivity to start date

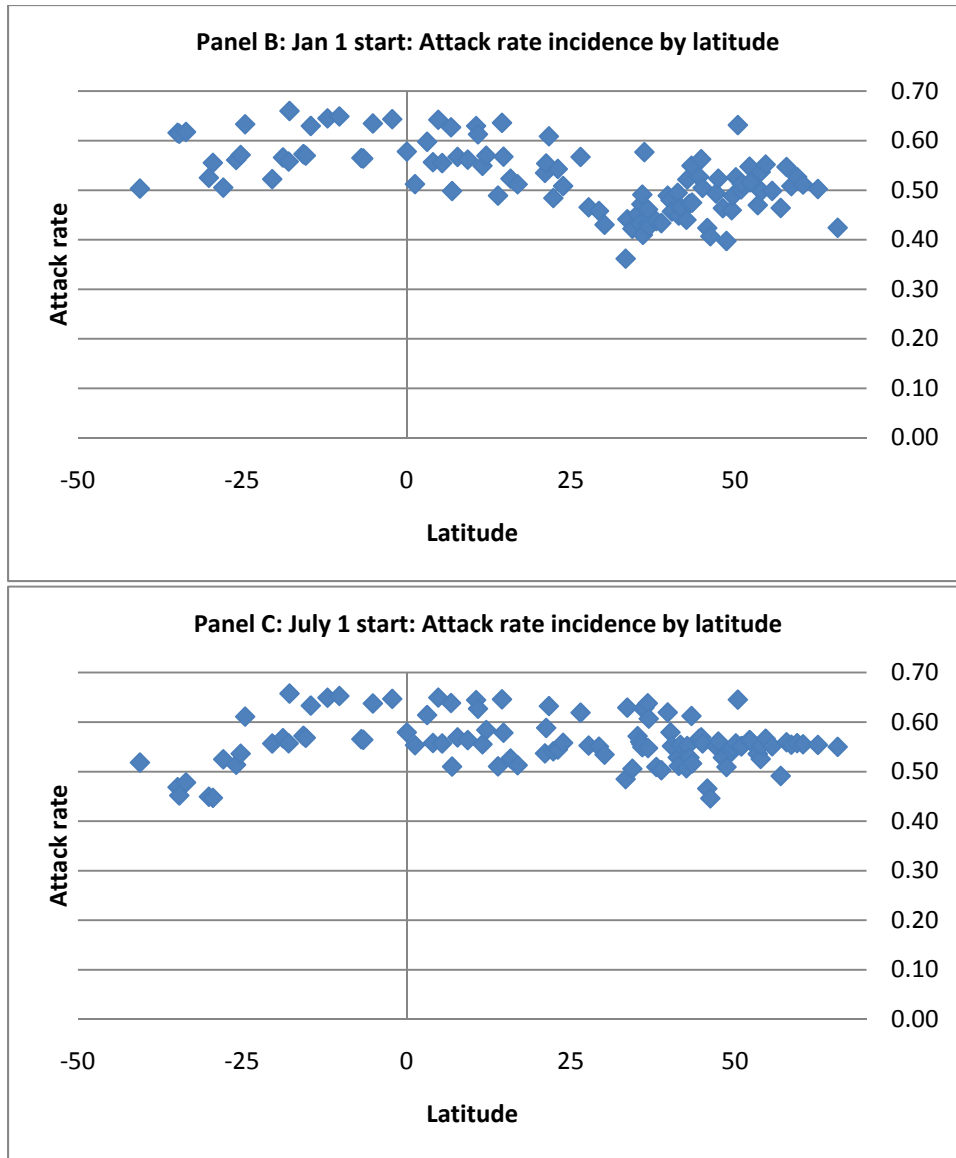


The incidence of attack rates across income groups is driven primarily by the combination of the seasonality assumptions and the correlation between income and latitude. Countries more distant from the equator have lower average seasonality factors, and so lower attack rates; and, higher income countries lie farther from the equator. Figure 10 displays the distribution of attack rates (after 3 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 in each case (the seasonality factor is at a minimum in summer). Variation in attack rates not due to latitude (the upward slope above ~mid 30s latitude in Panel B for example) is largely due to the structure of the international airline network and to variation in population density and urbanization rates across countries.

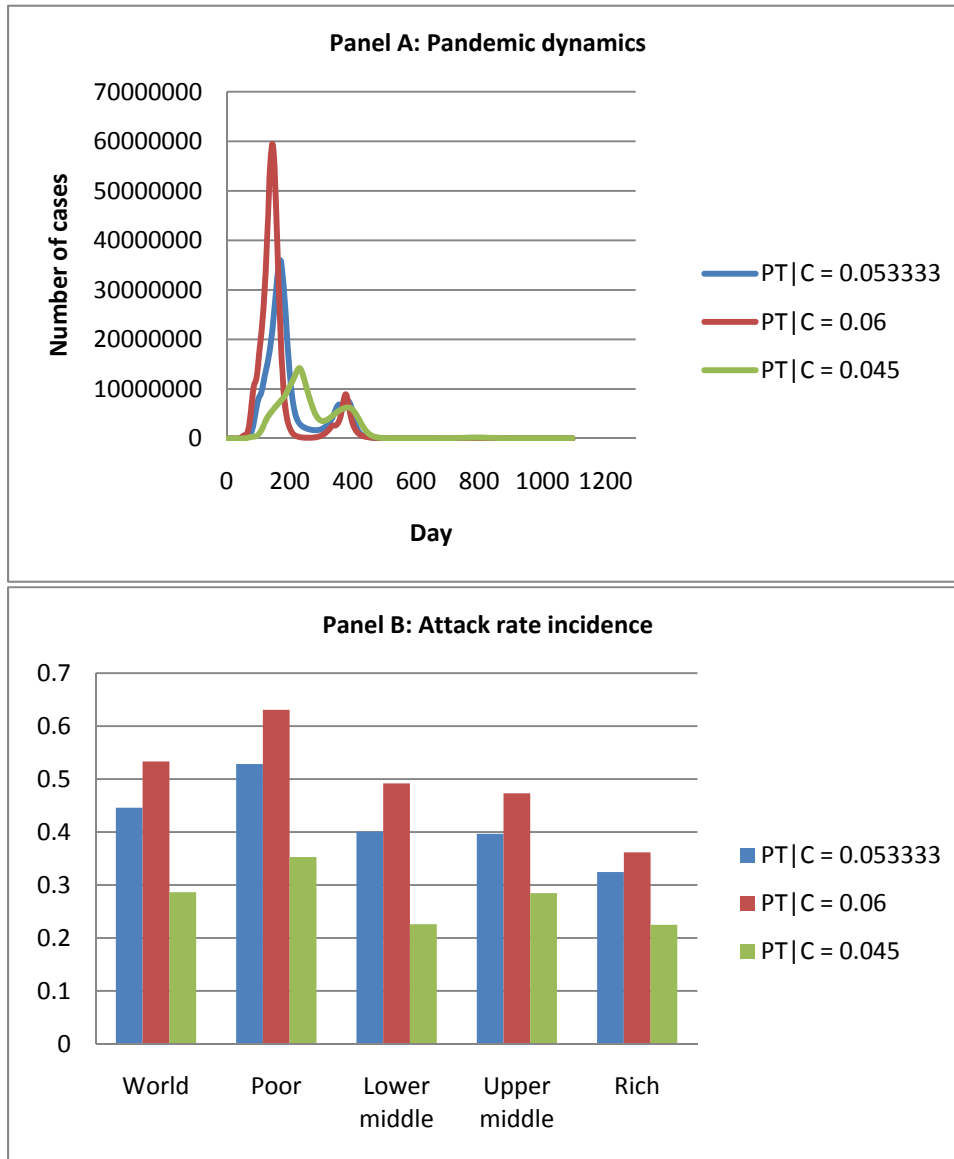
Figure 10: Latitude by income group and effect of latitude in baseline scenario





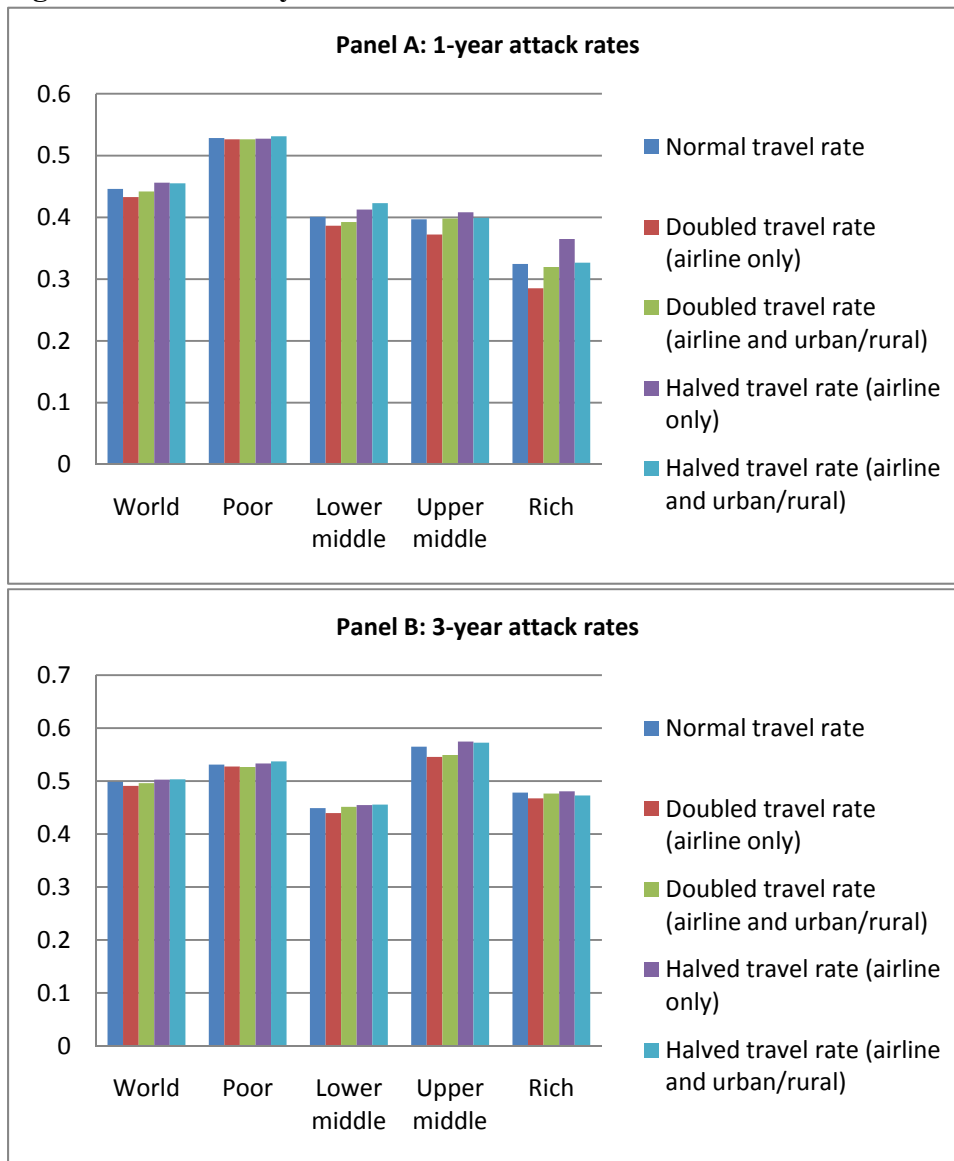
A higher value of $P(T|C)$ results in a faster pandemic and a higher attack rate (and the opposite for a lower $P(T|C)$ value) while retaining similar qualitative properties (see Figure 11). In a "severe" pandemic with $P(T|C) = 0.06$ the global 1-year attack rate rises to 53.3%, while in a "minor" pandemic with $P(T|C) = 0.045$ the global 1-year attack rate falls to 28.7%.

Figure 11: Baseline pandemic dynamics, sensitivity to $P(T|C)$



Results are generally insensitive to travel rate assumptions. Figure 12 shows that 1-year attack rates for wealthier countries vary slightly with a doubling or halving of travel rates, and that the impact on terminal attack rates after 3 years is minimal. Adjusting travel rates has two effects; a higher travel rate means a slightly faster pandemic spread (which might be expected to lead to higher attack rates), but a higher attack rate also shifts the season in which the pandemic peak occurs (which may increase or decrease attack rates depending on whether it shifts the pandemic peak towards or away from summer in the Northern Hemisphere). It appears that the second effect dominates; 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 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 each scenario we allocate stockpiles to each region and assume that the antivirals are distributed until the stockpile runs out. Antiviral policies are thus imposed on the model. In all scenarios, rich and upper middle income countries have stockpiles. Scenarios vary in the size of stockpiles distributed to lower middle income and poor countries. 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 (AVs) held by region r are used to treat people in all cities (including the rural area) within that region. We assume that AV 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 there is no AV remaining in the region.²² AVs are distributed to a proportion p^* of symptomatic infectious people.²³ In most runs, $p^* = 0.5$, implying that 1/3 of total new infectious (including asymptomatic cases) are treated each day, provided 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: it reduces the probability of transmission given contact by a multiplicative factor $(1 - e)$ where e is the efficacy of the antivirus, and it reduces the average duration spent infectious by one day, by uniformly shifting the probability distribution downwards. This means that an Infectious individual treated with AV will recover 4-7 days after becoming Exposed, with 30% recovering after 4 days, 40% after 5 days, 20% after 6 days, and 10% after 7 days. In most runs $e=0.6$.

We assume that there are no type 1 errors in treatment: AV doses are never used to treat non-infectious cases. While unrealistic, this assumption is relatively harmless, since we could increase the stockpile size in every city by some factor and receive the same qualitative results (treating x cases with no type 1 error is similar to treating $1.25x$ cases with a 20% false positive rate).

Choice of AV Scenarios

We focus on four AV stockpile scenarios described in the format a/b/c/d. This notation indicates that all regions in the Poor income group have a stockpile size of a% of their population, Lower middle income regions have a stockpile size of b%, and so forth.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 + a 4.2% stockpile for the outbreak country, Indonesia, which is in the Poor income group. The combined stockpile size of all rich regions is 91.5 and 25.8 million doses for all upper middle income regions (10% and 5%

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 outbreak source country.

²³We do not consider prophylactic use of AV. Although a Susceptible person treated with AV 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.

²⁴Although 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, since such a change may lead to additional people being treated in some regions but not in others.

of their respective populations) in all scenarios.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 1% of their population), while the movement from B to C is equivalent to rich countries paying for 28.4 million doses in Poor countries (equal to 1% of their population). In scenario D Lower Middle income countries have stockpiles equal to 1% of their populations and Indonesia has a stockpile of 9.2 million doses, which equals 4.2% of its population.

The impact of each scenario on attack rates at the end of one year depends on three factors: the infectiousness of the flu (i.e., $P(T|C)$), whether the flu starts on January 1 or July 1 and on the effectiveness of antivirals (e). We therefore present sensitivity analyses for all scenarios. We begin by describing the impact of each scenario on world attack rates, and then on the distribution of attack rates across country groups.

Effectiveness of Antivirals on World Attack Rates

Treatment with antivirus is highly effective in reducing pandemic attack rates. Table 3 shows the effect of AV treatment under scenarios A-C compared to the no-treatment scenario. The table also shows the impact of varying $P(T|C)$, the starting date of the pandemic and AV 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%), 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 no-AV baseline. Scenario A leads to an average of 2.85 fewer cases per dose, 7.16 fewer cases per dose for Scenario B, and 8.60 fewer cases per dose for Scenario C.

Table 3: Effect of AV on global 1-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

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% of their population size. Since 1/3 of infectious cases are treated, a region will not exhaust a 10% stockpile any time its attack rate is less than 0.3.

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 a/b/c/d refer to Poor/LowerMiddle/UpperMiddle/Rich regions. Low $P(T|C) = 0.045$, Medium $P(T|C) = 0.05333$, High $P(T|C) = 0.06$.

Table 4: Reduction in global influenza cases from AV, sensitivity to virulence

Scenario	$P(T C) = 0.045$			$P(T C) = 0.05333$			$P(T C) = 0.06$		
	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 no-AV baseline (so a positive value indicates fewer people infected). All figures are for the point 1 year after the pandemic commences.

Table 5: Reduction in global influenza cases from AV, sensitivity to AV efficacy

Scenario	$e = 0.6$			$e = 0.5$		
	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 no-AV baseline (so a positive value indicates fewer people infected). All figures are for the point 1 year after the pandemic commences.

Table 6: Reduction in global influenza cases from AV, sensitivity to start date

Scenario	January 1 start			July 1 start		
	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

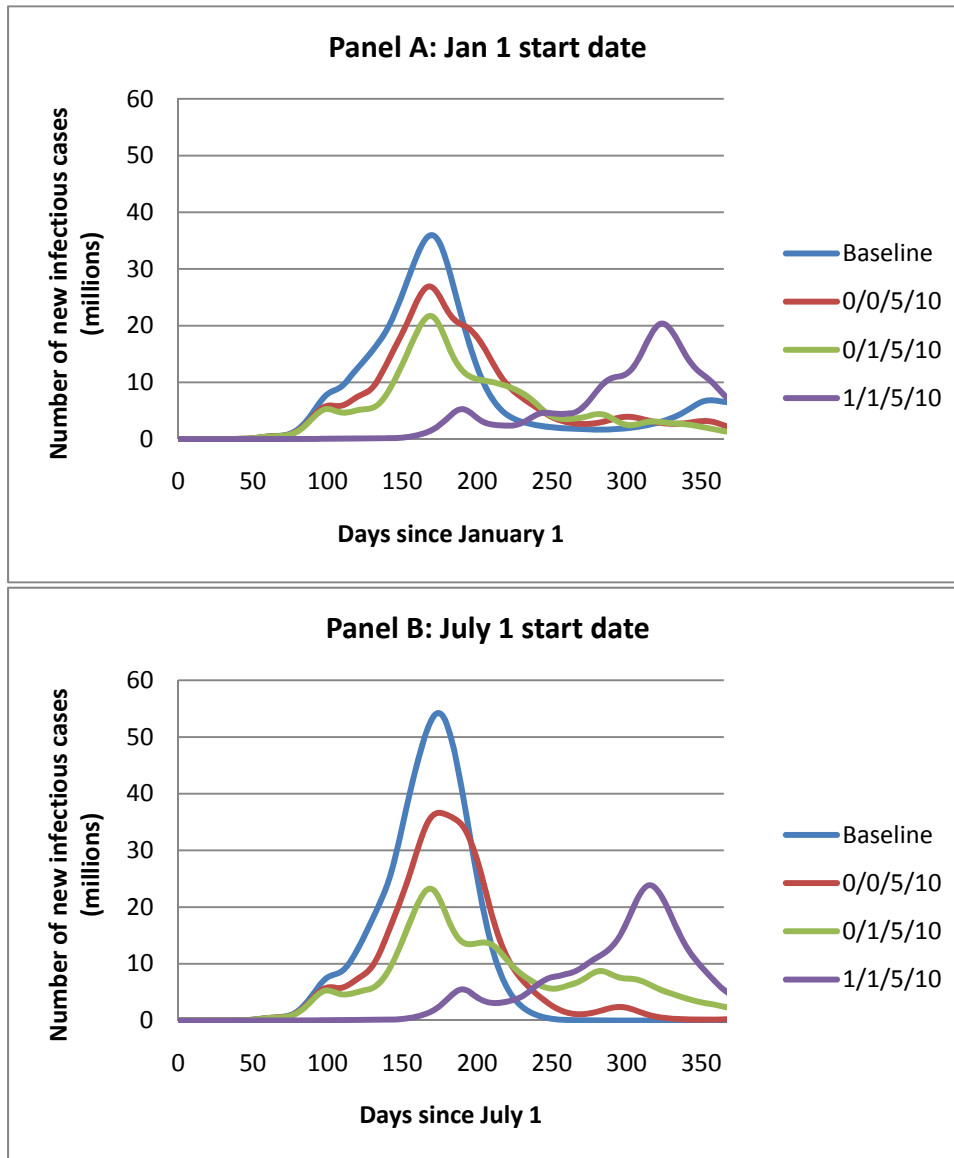
Note: Change in attack rate and number of cases refer to the reduction relative to the no-AV baseline (so a positive value indicates fewer people infected). All figures are for the point 1 year after the pandemic commences.

Increasing the infectiousness of the flu ($P(T|C) = 0.06$) reduces the effectiveness of AVs, while a lower value of $P(T|C)$ makes containment of the flu possible under scenario C. Scenario C allocates AV to all poor 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, and prevent a pandemic. But containment is not possible in other scenarios. Reducing the effectiveness of the AV dose from $e=0.6$ to $e=0.5$ reduces the number cases reduced per dose of AV 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 impact of AVs on the global attack rate is also affected by the start date of the pandemic. AVs are more effective in reducing attack rates, in both absolute and in percentage terms, if the pandemic starts on July 1. Note that most of the world's population lives in the Northern Hemisphere. Our seasonality assumptions imply that infectiousness in non-tropical regions is highest in winter and lowest in summer. This means that AV policies that slow the pandemic will tend to increase the average seasonality factor for a January 1 start date (which peaks in June, when the Northern Hemisphere seasonality factor is near minimum), and decrease the average seasonality factor for a July start date (which peaks in December, when the Northern Hemisphere seasonality factor is near maximum). Thus, the net impact of an AV policy is caused by both its direct effect from reducing infectiousness, and its indirect effect via the change in seasonality factor.

Figure 13 shows how AV treatment delays the pandemic (shifting the distribution rightwards). With a January 1 start date, delay blunts the effectiveness of AV in reducing the attack rate, because it pushes the pandemic peak towards Northern Hemisphere winter. With a July 1 start date, delay augments the effectiveness of AV in reducing the attack rate, because it pushes the pandemic peak away from Northern Hemisphere winter.

Figure 13: GEM pandemic dynamics with AV treatment



Effect of Antiviral Usage in One Region on Other Regions

In one and two-city SIR models, most of the gains from AV treatment accrued to the region in which the treatment was taken, but there were some positive spillover effects. These results also hold in the full GEM. Table 7 shows 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 AV to the 0/0/5/10 scenario in the base case, there is a large reduction in attack rate in Rich regions and in Upper Middle income regions, and negligible impacts in Lower Middle and Poor income regions.²⁶ Adding a stockpile in Lower Middle income regions

²⁶ In fact, there is a very small increase in attack rate in Poor and Lower Middle income regions. This is due to interaction with seasonality effects: AV in wealthier countries delays the pandemic slightly, which means that Northern Hemisphere countries (including India and China) have slightly higher average seasonality factors.

significantly reduces their attack rate while causing a small reduction in Rich country attack rates but a larger reduction in attack rates in Upper Middle income countries.²⁷ Moving to 1/1/5/10 adds an AV stockpile in poor countries, leading to a large attack rate reduction in poor countries and moderate benefits elsewhere.

Table 7: Impact of AV on 1-year attack rate, incidence by income group, standard parameters

AV scenario	#AV doses	World	Poor	Lower middle	Upper middle	Rich
Jan 1 start date, 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
Jan 1 start date, 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
Jan 1 start date, 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)						
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 a January 1 start date, and AV efficacy 0.6

The small external benefits to Rich countries of Lower Middle and Poor countries having stockpiles still holds when the flu starts on July 1, assuming it is moderately transmissible. A highly transmissible flu, however, sharply reduces the benefits to Rich countries of stockpiles in Lower Middle income and Poor countries, because a small stockpile is rapidly exhausted under a virus with a high reproductive rate. Indeed, in the January 1, High Transmissibility case Rich countries are slightly worse off in Scenario C compared to Scenario B. The fact that Lower Middle income countries themselves are worse off when using AVs (118 million additional cases relative to no treatment) in the High transmissibility scenario with the January 1 start date

²⁷ The latter effect is due primarily to a reduction in attack rates in Russia (which constitutes a sizeable proportion of the Upper Middle income population), which is at high latitude and so gains a large amount from delaying the pandemic peak past winter, given our seasonality assumptions.

follows largely from a higher attack rate in China, where the slight delay from a 1% stockpile pushes the pandemic peak towards winter (and so increases China's average seasonality factor).

It seems unlikely poor and lower middle income countries will purchase and maintain their own stockpiles of antivirus for use in the course of a pandemic. But the external benefits to rich countries suggest that there may be scope for a Pareto-improvement; rich countries may find it in their interest to pay for the costs of acquiring and distributing antivirals in poor countries. As discussed more fully below, Table 7 implies that the reduction in number of influenza cases in rich countries per AV dose administered in poor or lower middle income countries is 1/3 or more when the transmissibility of the flu is low to moderate.

An alternative is for rich countries to donate AV directly to the outbreak source. Consider a scenario in which rich 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 rich 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 $P(T|C)$ scenarios, because in these cases donation to the outbreak source reduces the R_0 below 1 and the pandemic is contained. The gains are also large with a July 1 start date, where delay is more valuable due to the interaction with seasonality. The external benefits of providing AV are lower when the AV efficacy is lower, which contributes to the negative impact of donation in the low AV efficacy scenario.

Table 8: Impact of AV on 1-year attack rate: Scenario D

P(T C)	Start date	AV efficacy	Attack rate for:	Policy: No donation	Policy: Donation to Indonesia	Difference
Low	1-Jan	0.6	World	0.178	0.001	0.176
Low	1-Jan	0.6	Rich	0.015	0.001	0.014
Medium	1-Jan	0.6	World	0.310	0.289	0.021
Medium	1-Jan	0.6	Rich	0.052	0.047	0.005
Medium	1-Jan	0.5	World	0.310	0.319	-0.009
Medium	1-Jan	0.5	Rich	0.052	0.068	-0.016
Medium	1-Jul	0.6	World	0.450	0.328	0.122
Medium	1-Jul	0.6	Rich	0.111	0.081	0.031
High	1-Jan	0.6	World	0.489	0.468	0.021
High	1-Jan	0.6	Rich	0.103	0.108	-0.005

²⁸ Nearly all rich regions do not consume their entire 10% stockpile (assuming Moderate $P(T|C)$) at all or within 1 year, and so the net effect of this policy is almost identical to one where rich countries simply paid for the additional doses in Indonesia.

Which strategy is more effective for rich countries: paying for AV doses to be divided among poor countries in proportion to population, or paying for AV doses to be targeted to the outbreak country? Under what circumstances will welfare in rich countries be increased by providing AVs to low income countries? Table 9 compares the effectiveness of purchasing doses for low 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 rich countries for a high virulence pandemic ($P(T|C) = 0.06$). But in all other cases, targeting the outbreak source is dramatically more cost effective than spreading doses throughout Poor or Lower middle income regions.

Table 9: Number of Rich country cases reduced per dose purchased for low income countries

Scenario transition:	P(T C)	Start date	Rich country cases reduced per extra AV dose purchased
Scenario A ->Scenario B	Low	1-Jan	0.74
Scenario A ->Scenario B	Medium	1-Jan	0.31
Scenario A ->Scenario B	Medium	1-Jul	0.33
Scenario A ->Scenario B	High	1-Jan	0.00
Scenario B ->Scenario C	Low	1-Jan	0.29
Scenario B ->Scenario C	Medium	1-Jan	0.22
Scenario B ->Scenario C	Medium	1-Jul	0.83
Scenario B ->Scenario C	High	1-Jan	-0.17
Scenario B ->Scenario D	Low	1-Jan	1.42
Scenario B ->Scenario D	Medium	1-Jan	0.52
Scenario B ->Scenario D	Medium	1-Jul	3.06
Scenario B ->Scenario D	High	1-Jan	-0.51

Table 9 suggests that providing AVs to developing countries in the event of a pandemic may pass a benefit-cost test. Although the percentage reduction in cases from providing AVs is small, millions of cases of the flu in rich countries would thereby be avoided, at a cost of 3-4 doses of antivirals per case avoided. 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 U.S. at \$187 per person, based on a 50% attack rate. Keogh-Brown et al. (2010) suggest that a mild influenza pandemic in the UK (similar to the 1957 or 1968 pandemics) would reduce GDP by 0.58% over the course of year; a more severe pandemic (with a case fatality rate of 1%) would reduce GDP by 4.5% over the course of a year. Typical estimates of the value of a statistical life in rich countries are in the millions of dollars, and so even a very low but positive case fatality rate would lead to a large value from reducing cases.

Sensitivity of Results to Proportion Treated in Poor Countries

The effectiveness of using AV to mitigate pandemics is a function of the effectiveness of the public health infrastructure in each country, and its ability to rapidly and accurately identify infectious individuals and provide them with AV treatment. In the primary simulation runs, we assume that 50% of symptomatic cases could be reached within a day, in any country that has a stockpile. These assumptions are arguably too optimistic, particularly for developing countries, though they are more pessimistic than the assumptions used in Colizza et al. (2007). By weakening these assumptions and reducing the proportion of symptomatic infectious cases who receive treatment, we examine the sensitivity to these assumptions.

Table 10 presents simulations based on a weak health infrastructure scenario in which only 30% of symptomatic infectious cases can be treated in poor and lower middle income countries, 40% in upper middle income countries, and (as before) 50% in rich countries. The effect of reducing the proportion who are treated is dramatic, both on the direct value of antiviral treatment in countries that have a stockpile, and on the external benefits to rich countries from providing antiviral doses to poorer countries. In terms of direct effects, the marginal benefit (in terms of world cases reduced) from an AV stockpile in upper middle and rich countries is reduced from 5.2 percentage points to 4.8 percentage points. The marginal benefit from adding a 1% 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% stockpile to poor countries is reduced from 5.6 percentage points to 1.3 percentage points.

Table 10: Attack rate after 1 year, standard vs. reduced proportion of infectious treated

		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
Poor	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
Rich	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% of new symptomatic infectious cases are treated (in regions that have an AV stockpile); weak health infrastructure means that 30% of cases are treated in poor and lower middle income regions, 40% in upper middle income regions, and 50% in rich regions. All scenarios have January 1 start date.

The benefits to rich countries from treatment in poor and lower middle income countries are also significantly reduced in the weak (compared to the strong) health infrastructure scenarios. The marginal reduction in the attack rate in rich countries from adding a 1% stockpile for lower middle income countries falls from 0.7 percentage points to 0.3 percentage points. Similarly for adding a 1% stockpile to poor countries, the benefit falls from 0.7 to 0.3. These results follow from the basic "increasing returns" property of the core SIR model; with a lower percentage of people treated, the effect of AV on reducing the reproductive rate is diminished proportionally, which has a greater than proportional impact on the benefits of AV from reducing attack rates.

These results suggest that if the weak health infrastructure parameters are a more accurate description of the real world, the benefits to rich countries of providing AVs to poor countries, based on self-interest, are greatly reduced, a result that is confirmed by Table 11. The results suggest, however, that investments that increase the number of infectious people than can be treated in poorer countries will have large benefits, and will be complementary to policies that provide antiviral doses to poor countries.

Table 11: Number of Rich country cases reduced per AV dose purchased in poor countries, sensitivity to weak health infrastructure

Scenario Transition:	P(T C)	Start date	Rich country cases reduced per extra AV dose purchased	
			Strong health infrastructure	Weak health infrastructure
Scenario A ->B	Medium	1-Jan	0.31	0.12
Scenario B ->C	Medium	1-Jan	0.22	0.09
Scenario B ->D	Medium	1-Jan	0.52	0.21

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 poor countries are unlikely to hold stockpiles of antiviral drugs, it may be in the self-interest of wealthy countries to (collectively) pay for purchase and distribute antivirals to poor countries, even without any altruistic or humanitarian motivations.

The payoff to providing antiviral doses to poor countries is illustrated by a simple, two-country SIR model. In that model there are significant externalities and complementarities in antiviral treatment: when one country treats more of its population this both reduces the attack rate in the other country and increases the marginal benefit from additional treatment in the other country.

It would, however, 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 Northern Hemisphere summer. And, the number and distribution of poor v. rich 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.

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

These results are, however, dependent on the transmissibility of the flu, its effectiveness in reducing infection and on the proportion of infectious who can realistically be identified and treated. Our simulations reveal that reducing the proportion of symptomatic infectious that can be treated from 50% (our base case assumption) to 40% in lower middle income countries and 30% in poor countries more than doubles the number of doses required to reduce a case of the flu in rich countries. Providing stockpiles to poor countries may still pass the “selfish” benefit-cost test. But, our results suggest that improving the delivery of health services in poor 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--277.
- Colizza, V., et al. (2007). Modeling the worldwide spread of pandemic influenza: Baseline case and containment interventions. *PLoS Med* 4(1), e13.
- Cooper, B. S., et al. (2006). Delaying the international spread of pandemic influenza. *Public Library of Science 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, 209-214.
- Ferguson, N. M., et al. (2006). Strategies for mitigating an influenza pandemic. *Nature*, advance online publication, 26 April 2006.
- Germann, T. C., et al. (2006). Mitigation strategies for pandemic influenza in the United States. *Proceedings of the National Academy of Sciences*. 103(15), 5935-5940.
- Glezen W. P., (1996). Emerging infections: pandemic influenza. *Epidemiology Review*, 18, 64–76.
- Grais, R. F., Ellis, J. H., and Glass, G. E. (2003). Assessing the impact of airline travel on the geographic spread of pandemic influenza. *European Journal of Epidemiology*, 18, 1065-1072.
- 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, 7794-7799.
- Hajdin, C. et al. (2009). Stochastic Equation-Based Model of a Global Epidemic (Version 3.0). RTI International, May 2009.
- Hethcote, H. W. and Yorke, J. A. (1984). Gonorrhoea transmission dynamics and control. *Lecture Notes in Biomathematics*, vol. 56. Berlin, Germany: Springer
- Hufnagel, L., Brockmann, D., and Geisel, T. (2004). Forecast and control of epidemics in a globalized world. *Proceedings of the National Academy of Sciences*. 101(42), 15124-15129.
- Keogh-Brown, M. R. et al. (2010). The Possible Macroeconomic Impact on the UK of an Influenza Pandemic. *Health Economics*. 19, 1345-1360.

Kermack, W. O. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proc. Roy. Soc. Lond. A* 115, 700-721.

Longini, I. M. Jr., et al., (2005). Containing pandemic influenza at the source. *Science*, 309, 1083-1087.

Lowen, A.C., et al. (2008). High temperature (30°C) blocks aerosol but not contact transmission of influenza virus. *Journal of Virology*, 82 (11), 5650-5652

Mossong, J. et al. (2008). Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLOS Medicine*. 5, 381-91.

Over, M., and Piot, P. (1993). HIV Infection and Sexually Transmitted Disease, in D.T. Jamison and others (eds.), *Disease Control Priorities in Developing Countries*. New York: Oxford University Press.

Rvachev, L. A. and Longini, Jr., I. M. (1985). A mathematical model for the global spread of influenza. *Mathematical Biosciences*. 75, 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), 226-233.

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, 881-889

Vynnycky, E. and Edmunds, W.J., (2008). *Epidemiol. Infect.*, 136, 166-179

Appendix A: Results for 2-city model, sensitivity to key parameters

β°	e	A	ρ_A	ρ_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

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 4x4 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), decompose this matrix into two parts: a 4x4 mixing matrix M_i and a 4x1 contact vector A_i , such that

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

where I_4 is the 4x4 identity matrix.

The vector A_i gives the absolute number of contacts per day for each age group; so for example $A_{\text{Jakarta}} = [7.29, 9.45, 6.41, 3.98]$ 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. Define the 4x1 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% of the population are in age group 0, 18.75% are in age group 1, 66.09% are in age group 2 and 5.52% are in age group 3, so the total contacts per day is the product $n_{\text{Jakarta}}' A_{\text{Jakarta}} = 6.94$.

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 from data on population density and urbanization rates. Mossong et al. measure physical contact rate matrices for eight European countries for 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 8 countries. The goal is to use these data to extract an underlying set of "core behavior" that can be assumed to hold in all countries, and will allow reconstruction of a C -matrix for every region in the GEM.

The Mossong data are used 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

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

day across the 8 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 GEM.

An M matrix for each region of the model by is estimated by assuming a structural form for the matrix, and then using the 8 Mossong matrices to estimate an underlying mixing matrix stripped of the country-specific n-vectors. Following Hethcote and Yorke (1984), define a 4x1 vector B_i :

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

and a 4x4 matrix G whose elements are defined by:

$$B_{ij}(1 - G_{ijk}) = M_{ijk}$$

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 jth 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 is what allows us to move away from random mixing, and captures the more complex structure where people can have differential mixing rates across age-groups.

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 That is, each age-group is assumed to have the same relative preference for mixing with its own age-group as for all other age-groups. Estimate the G-matrix that minimizes the sum of squared residuals across the 8 Mossong countries to find an “optimal” G matrix, G^* . M, and hence C matrices, for every region in the model are then reconstructed by substituting G^* and the region-specific age structures from real world population data into the equations above.

Contact rates are adjusted for population density, recognizing that areas with higher population density have higher contact rates. Every city and rural area is assigned to one of four categories; relative density $d=1$ for rural areas with residual urbanization <40%, $d=1.1$ for rural areas with residual urbanization >40%; $d = 1.3$ for low density urban areas (cities with density < 1,000 people per square kilometer); and $d=1.4$ for high density urban areas (for cities with density > 1,000 per square kilometer). In general, rural areas in high income countries have high

³⁰ This crucial assumption, made for simplicity, implies that there is no variation in total contacts per day 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.

³¹ Tested alternative forms showed that moving from a 16-unique-element G-matrix to a 4-unique element G-matrix had very little impact on the fit to the Mossong data, but that moving from a 4-element G-matrix to a 1-element G-matrix significantly worsened the fit.

residual urbanization while those in low income countries have low residual urbanization, and 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 GEM, by Income Group

Regions in the Poor 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,

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 Rich Income Group:

Australia, Austria, Belgium, Canada, Cyprus, Czech Republic, Denmark, Finland, France, Germany, Greece, Hong Kong S.A.R. 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 EFTA, Rest of North America