

Dynamic Health Policies for Controlling the Spread of Emerging Infections: Influenza as an Example

Reza Yaesoubi^{1,2*}, Ted Cohen^{1,2,3}

1 Center for Communicable Disease Dynamics, Harvard School of Public Health, Boston, Massachusetts, United States of America, **2** Division of Global Health Equity, Brigham and Women's Hospital, Boston, Massachusetts, United States of America, **3** Department of Epidemiology, Harvard School of Public Health, Boston, Massachusetts, United States of America

Abstract

The recent appearance and spread of novel infectious pathogens provide motivation for using models as tools to guide public health decision-making. Here we describe a modeling approach for developing dynamic health policies that allow for adaptive decision-making as new data become available during an epidemic. In contrast to static health policies which have generally been selected by comparing the performance of a limited number of pre-determined sequences of interventions within simulation or mathematical models, dynamic health policies produce “real-time” recommendations for the choice of the best current intervention based on the observable state of the epidemic. Using cumulative real-time data for disease spread coupled with current information about resource availability, these policies provide recommendations for interventions that optimally utilize available resources to preserve the overall health of the population. We illustrate the design and implementation of a dynamic health policy for the control of a novel strain of influenza, where we assume that two types of intervention may be available during the epidemic: (1) vaccines and antiviral drugs, and (2) transmission reducing measures, such as social distancing or mask use, that may be turned “on” or “off” repeatedly during the course of epidemic. In this example, the optimal dynamic health policy maximizes the overall population's health during the epidemic by specifying at any point of time, based on observable conditions, (1) the number of individuals to vaccinate if vaccines are available, and (2) whether the transmission-reducing intervention should be either employed or removed.

Citation: Yaesoubi R, Cohen T (2011) Dynamic Health Policies for Controlling the Spread of Emerging Infections: Influenza as an Example. PLoS ONE 6(9): e24043. doi:10.1371/journal.pone.0024043

Editor: Mike B. Gravenor, University of Swansea, United Kingdom

Received: April 4, 2011; **Accepted:** July 29, 2011; **Published:** September 6, 2011

Copyright: © 2011 Yaesoubi, Cohen. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: The work is supported by National Institutes of Health grants DP2OD006663 and U54GM088558. The content is solely the responsibility of the authors and does not necessarily represent the official views of the Office of the Director of the National Institute of Health and the National Institute of General Medical Sciences of the National Institutes of Health. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: reza.yaesoubi@gmail.com

Introduction

The recent appearance of novel human pathogens such as H1N1 and H5N1 influenza, and SARS has stimulated efforts to develop methods to determine health policies that allow for the rapid modification of public health interventions in response to varying epidemiological situations and changing availability of information [1,2]. In this paper, we examine the potential utility of *dynamic health policies* for controlling the spread of an emerging infectious disease. Dynamic health policies make *real-time* recommendations, in response to changing disease characteristics (e.g. infectivity, antimicrobial resistance levels), population characteristics (e.g. disease prevalence, proportion of individuals that are immune), and resource constraints (e.g. vaccines, antimicrobial drugs, personnel, and budget).

Most existing approaches for identifying optimal strategies for infectious disease control use simulation or mathematical models of disease spread to compare the performance of a *limited* number of *pre-determined* health policies. A number of these studies aim to identify optimal health policies for vaccine allocation *before* the start of an epidemic without explicitly considering interventions which can be employed during the epidemic [3–6]. A larger number of studies investigate the effect of both initial immunization and the

use of controlling interventions during epidemics, such as use of antiviral for treatment, case isolation, school closure, and internal travel and border restrictions [7–13].

Although these approaches can provide insight into which baseline strategies may best reduce the impact of epidemics, they are not generally structured to assist real-time decision making through the *dynamic* change of health recommendations as new data become available over the course of epidemic. In this paper, we focus on developing optimal dynamic health policies for controlling an emerging human pathogen. These policies allow decision-makers to use cumulative real-time data from the epidemic and current information about resource availability to guide their selection of possible interventions at any particular point in time.

We use a simplified model of influenza spread to illustrate the development and the employment of dynamic health policies. Control of influenza epidemics may involve: (1) reducing the susceptibility of uninfected individuals either before or during the epidemic (through vaccination or antiviral prophylaxis), (2) reducing contact rates in the population (through social distancing such as isolation of diagnosed cases, quarantine of households of diagnosed cases, closing of schools), and (3) reducing the infectiousness of infected individuals (through treatment or isolation).

While vaccines provide protection from seasonal influenza, the emergence of a novel strain would likely mean the absence (or shortage) of an effective vaccine for at least the first several months of the epidemic [14,15]. During this period, control strategies would largely rely on social distancing and potentially on stockpiles of antiviral drugs used for treatment and/or prophylaxis. Therefore, efforts to control an emerging influenza epidemic would be bounded by (1) the availability of effective vaccines and antiviral drugs, and/or (2) the availability of money and resources for vaccine procurement, diagnosis and treatment of new cases, and the actual implementation of transmission reducing interventions during the epidemic. Under such resource constraints, we define the *optimality* of a dynamic health policy as the efficient use of available resources (e.g. budget and medical system capacity) to maximize the overall health of the population (e.g. minimizing the number of deaths or hospitalizations, or maximizing other measures such as quality-adjusted life years).

In this paper, we define several broad types of interventions for controlling the spread of an emerging influenza virus. We use these crudely classified interventions and a simplified model of influenza transmission to illustrate the use of a dynamic optimization methodology (namely, Markov Decision Process [16]) to specify dynamic health policies. Here, rather than focusing on developing a comprehensive model of influenza spread, we demonstrate how these dynamic policies allow real-time decision making under different resource constraints. We also discuss how future research in this area can help make the implementation of these policies possible.

Method

In this section, we first discuss the set of possible interventions that can be used for controlling an influenza epidemic. We describe the effect of these interventions on disease spread as well as on the overall health-related and monetary consequences of the epidemic. We then define dynamic health policies for controlling the epidemic and characterize their optimality. We finally propose a methodology to identify the optimal dynamic health policies for controlling the epidemic.

A Model for Controlling the Spread of Influenza

Influenza Epidemic State and Decision Sets. An influenza epidemic is usually described by a SIR (Susceptible-Infective-Recovered) model, in which the individuals recovered from infection are assumed to acquire permanent immunity to that viral strain [7,17]. Let $X_S(t)$ denote the number of susceptibles, $X_I(t)$ denote the number of infectives, and $X_R(t)$ denote the number of recovered at time t . Since influenza epidemics usually last for several months and the number of deaths is generally small relative to population size, it is reasonable to assume that the population size does not change over the course of epidemic. For a population of a fixed size N , the state of the disease spread at any given time t can be identified by $s_t = (X_S(t), X_I(t))$. Let Ω denote the state space defined as $\Omega = \{(x_S, x_I) \in \mathbb{N}^2 \mid x_S + x_I \leq N\}$.

Decisions are made at points of time referred to as *decision epochs*. It is more convenient to assume that the decisions are made at discrete points of time, rather than continuously over time; hence, we assume that the set of decision epochs, \mathcal{T} , is discrete; that is $\mathcal{T} = \{t_0, t_1, t_2, \dots, T^H\}$, where T^H denotes the decision horizon length. We classify the possible interventions to control the spread of influenza into two categories: (1) “irreversible” interventions such as vaccination employed either before or during the epidemic which reduce the number of susceptibles, and (2) “reversible” interventions which can be turned on and off during the course of

epidemic to reduce the transmission of infection to susceptibles, such as hygienic interventions, social distancing, and treatment. We implement these two types of decisions in our model as follows:

1. **Vaccination:** At any decision epoch, conditional on the availability of vaccine, the decision maker will specify the number of susceptibles to vaccinate. If effective vaccines are abundant and vaccination has zero cost, this decision is trivial: vaccinate all susceptibles. However, throughout this paper we assume that vaccine, if available, is acquired at a price. We denote this decision by $z_t \in A_I$, where A_I is the set of possible values for the decision variable z_t . For cases where vaccine availability exceeds need, we may assume that $A_I = [0, N]$. For simplicity, we assume that vaccination at decision epoch t results in immunization by the next decision epoch. Hence, a decision to vaccinate z_t portion of susceptibles at time t results in the reduction of susceptibles to $X_S(t) - z_t$ by the next decision epoch.
2. **Transmission-reducing intervention:** These interventions may be either employed or lifted over the course of epidemic to reduce the transmission of infection to remaining susceptibles. These measures will include social distancing (e.g. school or public place closure), hygienic interventions (e.g. mask use), and treatment or isolation of cases. Let $A_T = \{0, 1, 2, \dots, M\}$ denote the set of such interventions, where 0 represents “no intervention.” We denote the transmission-reducing decision made at time t by $a_t \in A_T$.

We categorize “treatment” as a transmission-reducing intervention which can be turned on or off during the epidemic. Of course, when sufficient antivirals are available, all new cases receive appropriate treatment, in which case “treatment” will not be included as a decision in the model. However, under conditions of antiviral limitation, one may include “treatment” as an intervention which can be turned on or off during the epidemic; this situation may occur when the use of antiviral must be prioritized among population subgroups [3,4].

To control the epidemic, a policy maker will continue to make decisions until the prevalence of the disease is sufficiently low. The stochastic nature of transmission (which is especially important during emergence or eradication), prevents accurate identification of the time when the disease will be eradicated; hence, we consider an infinite decision horizon: $T = \{t_0, t_1, t_2, \dots\}$ [16]. Although the decision horizon is infinity, we assume the decision making process stops when there are no more infectious individuals in the population.

Rewards and Transition Probabilities. As the result of vaccination employed at time $t = \{0, 1, 2, \dots\}$, the number of susceptibles is reduced from $X_S(t)$ to $X_S(t) - z_t$ by the next decision epoch, and the policy maker receives a reward $r(s_t, z_t) = -pz_t$, where p is the unit price of vaccine.

As the result of a transmission-reducing intervention employed at time $t = \{1, 2, \dots\}$, i.e. $a_t \in A_T$, the spread of disease at the next decision epoch is determined by the probability distribution $P(\cdot \mid s_t, a_t)$, and the policy maker receives a reward $r(s_t, a_t)$.

The reward $r(s_t, a_t)$ can be characterized in several ways and the choice of reward structure should reflect the policy maker’s set of priorities. For example, if the policy maker wants to minimize the total number of individuals infected over the course of epidemic, then reward $r(s_t, a_t)$ can be simply defined as $r(s_t, a_t) = E[I(t) \mid s_t, a_t]$, where $I(t)$ is the number of new infections during the period t . However, efforts to control epidemics may be bounded by the availability of medical resources, such as vaccines, medical personnel, and antiviral drugs, and monetary resources for vaccine

procurement, diagnosis and treatment, and implementation of interventions. In these situations, where both health-related outcomes and the resource consumption level are essential for determining the optimality of a health policy, a more comprehensive reward function is needed. A common approach for defining optimality in these situations is to assume that the policy maker's objective is to maximize the population's net monetary benefit [18]. To characterize the reward $r(s_t, a_t)$ accordingly, we must define several additional parameters:

- λ : policy maker's willingness-to-pay (WTP) for health.
- c : cost incurred for each incident infection (this may include diagnosis, treatment and other indirect costs).
- $c_T(a_t)$: cost of implementing the intervention $a_t \in A_T$ for period t ; we assume $c_T(0) = 0$.
- $u(s_t, a_t)$: expected costs incurred during the period t if the disease spread at time t is at state s_t and the policy maker implements the intervention $a_t \in A_T$ at decision epoch t ; $u(s_t, a_t) = cE[I(t)|s_t, a_t] + c_T(a_t)$.
- w : loss in health (quantified by quality-adjusted life years) due to infections.
- $q(s_t, a_t)$: expected loss in population's health during the period t if the disease spread at time t is at state s_t and the policy maker uses intervention $a_t \in A_T$ at decision epoch t ; $q(s_t, a_t) = wE[I(t)|s_t, a_t]$.

Now, the reward $r(s_t, a_t)$, defined as the expected net monetary benefit during the period $[t, t + \Delta t]$ if the disease spread at time t is at state s_t and the policy maker chooses intervention $a_t \in A_T$ at decision epoch t , is calculated by $r(s_t, a_t) = -\lambda q(s_t, a_t) - u(s_t, a_t)$.

Decision Rules, Health Policies and Optimality. A *decision rule* prescribes an action for each state for a specified decision epoch. For decision epoch $t=0$, a decision rule is a function $d_0 : \Omega \rightarrow A_I$, which specifies the number $z_I \in A_I$ of susceptibles to vaccinate given the initial disease state $(X_S(0), X_I(0)) \in \Omega$. For decisions that must be made during the course of the epidemic ($t = \{1, 2, \dots\}$), we focus on a *Markovian* decision rule because this is the most convenient to implement and evaluate. Such decision rules are functions $d_t : \Omega \rightarrow A_T \times A_I$, which for each state of the disease spread, assign a transmission-reducing intervention $a_t \in A_T$ and a proportion $z_t \in A_I$ of susceptibles to vaccinate. This decision rule is said to be *Markovian* (memoryless) because it depends on previous disease states and previously utilized interventions only through the current state of the disease spread.

A *policy* specifies a decision rule to be used at all decision epochs. In other words, a policy π is a sequence of decision rules $\pi = (d_0, d_1, d_2, \dots)$. We call a policy *stationary* if $d_t = d$ for $t \in \{1, 2, \dots\}$; that is, the policy prescribes the same decision for the given state s regardless of the time period in which this state is reached. In this paper, we are only interested in characterizing stationary health policies since they are the most feasible to implement in practice and their optimality can be proven [16].

Assuming that the influenza spread is at state s_1 at decision epoch $t=1$, the expected total discounted reward induced by policy $\pi = (d_0, d, d, \dots)$ over the course of epidemic is calculated as:

$$v^\pi(s_1) = E^\pi \left[\sum_{t=1}^{\infty} \gamma^{t-1} r(s_t, d(s_t)) | s_1 \right], \tag{1}$$

where $\gamma \in [0, 1]$ is the discount factor to account for the time value of the future rewards.

Now, assuming that at decision epoch $t=0$, the state of the disease spread is s_0 , the expected total reward induced by policy $\pi = (d_0, d, d, \dots)$ is:

$$\begin{aligned} V^\pi(s_0) &= r(s_0, d_0(s_0)) + \gamma v^\pi(s_1) \\ &= -pd_0(s_0) + \gamma v^\pi(X_S(0) - d_0(s_0), X_I(1)). \end{aligned} \tag{2}$$

Let Π denote the set of all possible stationary policies $\pi = (d_0, d, d, \dots)$. We say that a policy $\pi^* \in \Pi$ is optimal whenever:

$$V^{\pi^*}(s_0) \geq V^\pi(s_0) \text{ for each } s_0 \in \Omega \text{ and all } \pi \in \Pi. \tag{3}$$

Define $v^*(s) \equiv v^{\pi^*}(s)$, where $v^{\pi^*}(s)$ is calculated by Eq. (1) for $\pi = \pi^*$.

Implicit in the definition of stationary policies is the assumption that model parameters are known and do not change over time. This assumption allows the existence of an optimal stationary policy $\pi = (d_0, d, d, \dots)$, where the function $d(\cdot)$ is time independent. Undoubtedly, this assumption may be violated for emergent influenza epidemics, and the model parameters may need to be updated as new data accrues over time. We return to this issue in later sections and discuss how stationary policies can be determined in this setting. But, for clarity of presentation, we first assume that the parameters of the influenza model are constant and estimable from the initial spread of the epidemic.

A Markov Decision Process Formulation for Influenza Spread

There are several dynamic optimization methodologies that can be used to find the policy π^* defined in inequality (3). The most appropriate optimization method depends on the complexity of the underlying epidemic model, the observability of the epidemic state, and the desired level of computational efficiency. As a rule of thumb, finding the *exact* optimal policy π^* becomes more challenging (and sometimes impossible) as (1) the complexity of epidemic model increases or (2) the uncertainty around the true state of the disease spread arises. Several methodologies can help identify the *approximate* optimal policy π^* even when the epidemic model is relatively complex and only some probabilistic knowledge about the current state of the epidemic can be obtained. We discuss the challenges of optimal decision-making under uncertainty at further length in the Discussion section.

If the disease dynamics can be modeled by a discrete-time Markov chain and the state of the epidemic is observable over the course of epidemic, then the stationary optimal health policy π^* in inequality (3) can be efficiently obtained through Markov Decision Process (MDP) [16]. In our illustrative example of a novel influenza epidemic, we make several simplifying assumptions about the spread of influenza to be able to use MDP. We enumerate these assumptions in the following subsections; we then discuss how these simplifying assumptions can be relaxed in further work.

We describe epidemic influenza with a SIR model, in which individuals acquire permanent immunity through infection or vaccination. We do not consider the possibility of changes in the population size due to birth, immigration, or death in order to simplify the analysis. Let Δt denote the time interval between two consecutive decision epochs. We assume that individuals become infected only through contact with other infectious members of the population, and that contacts during the interval $[t, t + \Delta t]$ occur according to a homogenous Poisson process, with rate $\mu \Delta t$.

We assume that a susceptible who is infected during period $[t - \Delta t, t]$ becomes infectious and symptomatic at time t and interacts with the rest of the population during period $[t, t + \Delta t]$; the individual is then removed from the population (or recovered) at time $t + \Delta t$. As we will describe later, these simplifying assumptions allow the use of Markov decision process to identify the optimal decision at each decision epoch. In the Discussion section, we explain how these assumptions can be relaxed.

Let τ denote probability that a susceptible person becomes infected upon contact with an infectious individual and $\beta(t)$ denote probability that the next interaction of a random susceptible person is with an infectious person. When social distancing has not been used and mixing is homogenous, $\beta(t)$ is equal to $\beta(t) = X_I(t)/N$. Variables τ and $\beta(t)$ can be respectively modified by “hygienic interventions” (reducing the chance of transmission given contact between infectious and susceptible individuals) and “social distancing” (reducing the likelihood of contact between susceptible and infectious individuals). Let $\varphi(t)$ denote overall probability that a susceptible person becomes infected. This probability is calculated in [19] as:

$$\varphi(t) = 1 - e^{-\mu\Delta t\beta(t)\tau}. \tag{4}$$

Hence, given the state of the disease spread, i.e. $s_t = (X_S(t), X_I(t))$, the number of new infections during period $[t, t + \Delta t]$, denoted by $I(t)$, will have a binomial distribution with number of trials $X_S(t)$ and the probability of success $\varphi(t)$:

$$P_{I(t)}(i | X_S(t), X_I(t)) = \begin{cases} \binom{X_S(t)}{i} \varphi(t)^i (1 - \varphi(t))^{X_S(t) - i}, & \text{for } 0 \leq i \leq X_S(t), \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

The transition probabilities of the Markov chain $\{(X_S(t), X_I(t)) : t = 0, 1, \dots\}$ can then be calculated by (refer to Text S1 for detailed steps):

$$\Pr\{(X_S(t + \Delta t), X_I(t + \Delta t)) = (x_S, x_I) | X_S(t), X_I(t)\} = \begin{cases} P_{I(t)}(x_I | X_S(t), X_I(t)), & \text{for } x_S + x_I = X_S(t), 0 \leq x_S \leq X_S(t), 0 \leq x_I \leq X_S(t), \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

We consider two types of interventions: (1) vaccination and (2) transmission-reducing interventions. First, we assume that no vaccines will become available any time during the epidemic; therefore, the health policy $\pi = (d, d, \dots)$ only specifies the optimal transmission-reducing intervention $a_t^* \in \mathcal{A}_T$ to implement at decision epoch $t \in \{1, 2, \dots\}$. We will show later how optimal vaccination decisions can be made once vaccine becomes available. Also, for simplicity, we assume that only one transmission-reducing intervention is available; hence $\mathcal{A}_T = \{0, 1\}$. Selecting intervention 1 at decision epoch t , i.e. $a_t = 1$, reduces probability $\varphi(t)$ in Eq. (4) to $\varphi(t) = 1 - e^{-(1 - \alpha_T)\mu\Delta t\beta(t)\tau}$, where $\alpha_T \in [0, 1]$ is the fractional reduction in the infection transmission rate.

Let $P(s' | s, a)$ denote the probability that the influenza epidemic will be in state s' at decision epoch $t + \Delta t$, given that the state is s at decision epoch t and the policy maker implements intervention $a \in \{0, 1\}$ at decision epoch t . This probability is calculated by transition probability (6).

Since we first assume that no vaccine is available during the course of the epidemic (we relax this assumption below), the policy maker can only select decisions from set $\mathcal{A}_T = \{0, 1\}$. Given this assumption, the optimal solution to the policy maker’s problem during the course of epidemic (i.e. maximizing Eq. (1)) is obtained by solving the following set of recursive equations [16]:

$$v^*(s) = \max_{a \in \{0, 1\}} \{r(s, a) + \gamma \sum_{s' \in \Omega} P(s' | s, a) v^*(s')\}, \text{ for } s \in \Omega. \tag{7}$$

By definition, function $v^*(s)$ returns the expected total discounted reward induced by the optimal policy over the remaining course of epidemic if the current state of epidemic is s . Therefore, having found the function $v^*(s), s \in \Omega$, by solving the set of equations (7), we can then determine the optimal transmission-reducing decision for a given state s by:

$$a^*(s) = \operatorname{argmax}_{a \in \{0, 1\}} \{r(s, a) + \gamma \sum_{s' \in \Omega} P(s' | s, a) v^*(s')\}. \tag{8}$$

Results

To illustrate the use of the proposed methodology, we consider the case of an influenza outbreak in an English boarding school reported in [20] and recently used by [2] and [21]. The population consisted of $N = 763$ students and the infection was believed to be introduced by one student returning from Asia. The situation satisfies many requirements of a simple SIR model, particularly since no specific intervention was employed during the outbreak.

For a population of size N , the transition probability matrix of the Markov chain $\{(X_S(t), X_I(t)) : t = 0, 1, \dots\}$ is of size $0.5(N + 1)^2$, which causes computational problems for our effort to identify optimal health policies. To overcome this computational difficulty, several effective alternative solutions have been proposed in the literature of dynamic optimization; the reader is referred to [22] for comprehensive discussion. One approach to reduce the state space of the Markov chain $\{(X_S(t), X_I(t)) : t = 0, 1, \dots\}$ is “state aggregation”, in which the Markov chain $\{(X_S(t), X_I(t)) : t = 0, 1, \dots\}$ is approximated by the Markov chain $\{(\Theta_S(t), \Theta_I(t)) : t = 0, 1, \dots\}$, where $\Theta_C(t), C \in \{S, I\}$, is the proportion of population in class $C \in \{S, I\}$ at time t . Detailed steps for how one can make these approximations are provided in the attached Text S1. Note that although we consider a relatively small population here, however, the approximation method briefly described above can also be used for larger populations.

The influenza spread model described previously has two parameters, μ and τ , which should be estimated from the data. Using maximum likelihood estimation, we estimate $\mu = 20$ per day and $\tau = 13.7\%$. Figure 1 shows each day’s expected number of new infections calculated by the model versus the observed data. Note that in Figure 1, the observations on days 1 and 2 are not shown. The reason is as follows: since in the approximate Markov chain $\{(\Theta_S(t), \Theta_I(t)) : t = 0, 1, \dots\}$, the states of influenza spread are aggregated, the model is not able to accurately capture the state of the epidemic when the number of infections is very low. For the observations presented in Figure 1, the number of new cases in days 1 and 2 are respectively 1 and 6, which are too low to be captured by the approximate Markov chain $\{(\Theta_S(t), \Theta_I(t)) : t = 0, 1, \dots\}$. For detailed discussion, refer to Text S1 and [19].

To determine optimal dynamic health policies for this population, we use the following arbitrary settings. We consider one transmission-reducing intervention which reduces the rate of

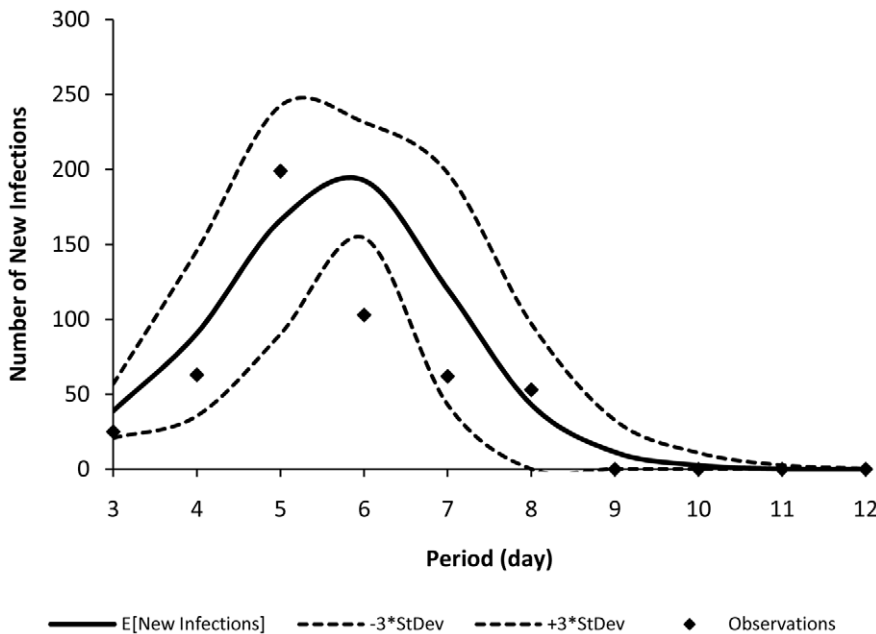


Figure 1. Observed new cases in the English boarding school versus the model’s predictions. This figure shows each day’s expected number of new infections calculated by the model (presented by solid curve) versus the observed data (presented by dots). The dotted curves show the model’s expected number of new infections ± 3 times its standard deviation.
doi:10.1371/journal.pone.0024043.g001

infection transmission by 20%, i.e. $\alpha_T=0.2$, and costs $c_T = \$,2000$ per day. One such transmission-reducing intervention might be “having all students wash their hands twice a day”. The vaccine price is set to $p = \$25$; each incident infection costs $c = \$100$ to diagnose and treat and results in a health loss of 0.00342 QALY (assuming that the treatment period lasts 5 days during which the health quality of the patient is reduced by 25%: $5 \times 0.25/365 = 0.00342$).

Optimal policies during the epidemic are obtained by solving the set of equations (7) for $v^*(\cdot)$ and then using Eq. (8) to find the optimal decision $a^*(s)$ for each state $s \in \Omega$. We use a policy iteration algorithm [16] with a discount factor $\gamma = 0.99$ to solve the set of equations (7).

Figure 2 displays the optimal health policies for WTP for health equal to $\$,25000/QALY$. The conditions for recommending the use of the transmission-reducing intervention are presented as the grey regions within the triangle in the upper part of each figure. For example, the transmission-reducing intervention should be implemented if the state of the disease is $(\theta_S, \theta_I) = (0.5, 0.2)$, and should not be used when the state of the disease is $(\theta_S, \theta_I) = (0.15, 0.4)$. Text S1 includes additional health policies corresponding to different levels of WTP.

As discussed previously, we must also identify the optimal number of susceptibles to vaccinate when an effective vaccine becomes available. Let us assume that the entire population is susceptible before the start of epidemic, i.e. $X_S(0) = x_S^0 = N$, among which x_I^1 susceptibles become infectious by time $t = 1$. If we assume that effective vaccines become available at time $t = 0$ at a market price p , then the policy maker determines the number the susceptibles to immunize, z_0 , by solving the following optimization problem:

$$V(x_S^0) = \max_{0 \leq z_0 \leq x_S^0 - x_I^1} v^*(x_S^0 - z_0 - x_I^1, x_I^1) - pz_0, \quad (9)$$

in which, the variable z_0 only takes integer values over the interval $[0, x_S^0 - x_I^1]$.

The rectangle in bottom of Figure 2, labeled as “% of susceptibles to immunize at $t = 0$,” represents the optimal health policies for immunization before the start of epidemic; again, if vaccine is not yet available, it may not be possible to achieve these levels of immunization. These recommendations are obtained by solving problem (9).

Employing Dynamic Health Policies to Control Influenza Spread

In this section, we discuss how the optimal dynamic health policy for using the transmission-reducing intervention and vaccinating additional susceptible individuals in decision epochs is determined as an epidemic progresses and new data become available.

Table 1 shows the observable information that accrues during an epidemic. Before the start of epidemic, at time $t = 0$, the policy maker obtains an estimate for $X_S(0)$, denoted by \hat{x}_S^0 . For a novel strain of influenza, we assume that the entire population is susceptible, hence $\hat{x}_S^0 = N$ for these situations. The policy maker also obtains a prior distribution for the number of susceptibles who may become infectious by time $t = 1$. Let the number of infectives at time $t = 1$ be randomly distributed according to probability mass function $P_{X_I(1)}(\cdot)$ with support $[0, \hat{x}_S^0]$. The policy maker can then use the following optimization problem to find the number of susceptibles to immunize at time $t = 0$:

$$z_0^* = \arg \max_{0 \leq z_0 \leq \hat{x}_S^0} \sum_{x_I=0}^{\hat{x}_S^0 - z_0} v^*(\hat{x}_S^0 - z_0 - x_I, x_I) P_{X_I(1)}(x_I) - pz_0, \quad (10)$$

in which, the variable z_0 only takes integer values over the interval $[0, \hat{x}_S^0]$.

Using observation of the number of new infections occurring during the epidemic, we update our knowledge on the state of the epidemic for the next decision epochs as follows. At the beginning of

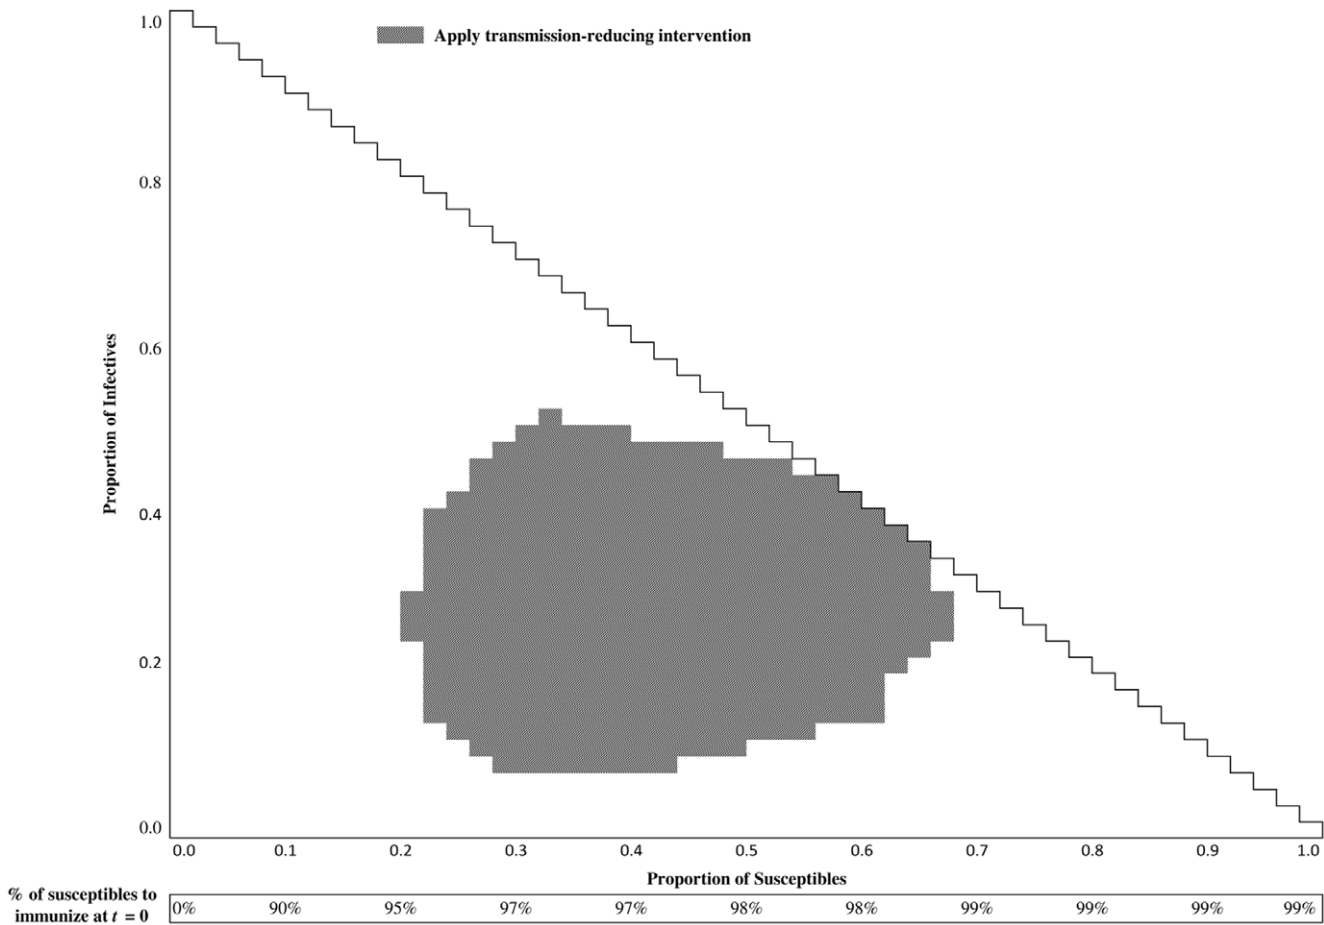


Figure 2. Optimal stationary health policy for $\lambda = \$,25000/QALY$. The conditions for recommending the use of the transmission-reducing intervention is presented within the triangle in the upper part of each figure. In disease states consistent with those that are captured within the grey cells, the intervention should be used. For instance, the transmission-reducing intervention should be implemented if the state of the disease is $(\theta_S, \theta_I) = (0.5, 0.2)$, and should not be used when the state of the disease is $(\theta_S, \theta_I) = (0.15, 0.4)$. The rectangle in bottom of this figure, labeled as “% of susceptibles to immunize at $t = 0$,” represents the optimal health policies for immunization before the start of epidemic given that vaccine is available at a market price.
doi:10.1371/journal.pone.0024043.g002

period $t = 1$, \hat{x}_I^1 new infections are observed which implies that the number of susceptibles at this time is $\hat{x}_S^1 = \hat{x}_S^0 - z_0^* - \hat{x}_I^1$. At the end of each period t , \hat{I}_t new infections are observed. This observation implies that (1) the number of susceptibles at the beginning of period $t + 1$ is $\hat{x}_S^{t+1} = \hat{x}_S^t - \hat{I}_t$ and (2) the number of infectives at time $t + 1$ is \hat{I}_{t+1} , since we assume that all infectives at time t are removed or recovered by the next decision epoch (see Table 1).

Now, knowing that the epidemic is at state $(\hat{x}_S^t, \hat{x}_I^t)$ and assuming that no vaccine is available at time t , the optimal

transmission-reducing decisions for period t is obtained by solving Eq. (8) (which can also be summarized in form of the policy shown in Figure 2).

Now assume that at decision epoch t , vaccines for the epidemic strain of influenza either become available for the first time or that depleted stocks of vaccine have now been replenished. Let $\hat{v}(x_S, x_I; a_t, z_t)$ denote the expected reward if influenza spread is in state (x_S, x_I) at time t , the transmission-reducing intervention $a_t \in \{0, 1\}$ is implemented during period $[t, t + \Delta t]$, and $z_t \leq x_S$

Table 1. Observed information over the course of an epidemic.

Period	0	1	2	...	t
Observed number of susceptibles, $\hat{X}_S(t)$	\hat{x}_S^0	$\hat{x}_S^1 = \hat{x}_S^0 - z_0^* - \hat{x}_I^1$	$\hat{x}_S^2 = \hat{x}_S^1 - \hat{I}_1$		$\hat{x}_S^t = \hat{x}_S^{t-1} - \hat{I}_{t-1}$
Observed number of infectives, $\hat{X}_I(t)$		\hat{x}_I^1	$\hat{x}_I^2 = \hat{I}_1$		$\hat{x}_I^t = \hat{I}_{t-1}$
Observed new infections, $\hat{I}(t)$		\hat{I}_1	\hat{I}_2		\hat{I}_t

doi:10.1371/journal.pone.0024043.t001

susceptibles are vaccinated by the decision epoch $t + \Delta t$. Then, function $\hat{v}(x_S, x_I; a_t, z_t)$ is calculated as:

$$\hat{v}(x_S, x_I; a_t, z_t) = r(x_S, x_I; a_t) - pz_t + \gamma \sum_{s' \in \Omega(z_t)} P(s'|x_S, x_I; a_t) v^*(s') \tag{11}$$

where $\Omega(z_t) = \{(x'_S, x'_I) \in \Omega | x'_S = x_S - z_t\}$. The optimal transmission-reducing intervention (a_t^*) and the number of susceptibles to vaccinate (z_t^*) during the period $[t, t + \Delta t]$ is then determined by solving the following problem:

$$\begin{aligned} & \max \quad \hat{v}(\hat{x}'_S, \hat{x}'_I; a_t, z_t), \\ & 0 \leq z_t \leq \hat{x}'_S, \\ & a_t \in \{0, 1\} \end{aligned} \tag{12}$$

where $\hat{v}(x_S, x_I; a_t, z_t)$ is given by Eq. (11) and the variable z_t only takes integer values over the interval $[0, \hat{x}'_S]$.

As briefly pointed out before, to determine the stationary health policy $\pi^* = [a^*(\cdot)]$, we assumed that the parameters of the influenza spread model are all known and do not change over time. It is, however, more realistic to assume that as new data become available over the course of epidemics, policy makers also update the parameters of the underlying transmission model. Table 2 outlines a procedure describing how dynamic health policies can be employed when policy maker is using the real-time data to also update the parameters of the influenza spread model.

Optimal Dynamic Health Policies under Resource Constraint

The optimal health policies developed in the previous section do not consider the availability of resources (e.g. vaccines, budget, and antiviral for treatment) while making health recommendations. For example, the policy presented in Figure 2, recommends vaccinating 99% of the remaining susceptibles when the initial proportion of susceptibles is 0.8; this recommendation does not account for the number of vaccines that are actually available. Similarly, the recommendations for turning on and off the transmission-reducing intervention do not account for the budget and resources necessary to initiate or terminate these recommendations. The framework proposed here can be expanded to incorporate different forms of resource constraints while con-

structing dynamic health policies. In this section, we discuss policy development under three types of resource limitations.

Let us first consider the case where n vaccines are available at time $t = 0$, and the policy maker must decide how many susceptibles to vaccinate using these n vaccines. For now, let us assume that no additional vaccines will become available during the epidemic and any unused vaccine by time $t = 1$ will be lost (e.g. shipped to other communities). The policy maker can now use Figure 3 to select the WTP for which a feasible vaccination recommendation can be implemented. For a given initial proportion of susceptibles θ_S^0 , Figure 3 identifies the number of vaccines required for each value of WTP (λ) through solving the optimization problem (10). As an example, for $\theta_S^0 = 0.5$, if the policy maker's sole objective is to minimize cost (i.e. $\lambda = \$0/\text{QALY}$) 352 vaccines are required and for any WTP $\lambda \geq \$2,5000/\text{QALY}$, 367 vaccines should be used. Note that for $\theta_S^0 = 0.5$, vaccinating less than 352 susceptibles results in a monetary loss, and vaccinating more than 367 susceptibles does not increase the population's expected net monetary benefit for any $\lambda \geq \$2,5000/\text{QALY}$.

Now to demonstrate how a vaccine supply constraint can be accounted for in constructing an optimal health policy, let us assume that for our population, at time $t = 0$ (before the start of epidemic), there are $n = 550$ vaccines available (the same approach can be followed for all other decision epochs). When $n = 550$ vaccines are available, then according to Figure 3, for any initial proportion of susceptibles $\theta_S^0 < 0.76$ any health policy with $\lambda \geq \$2,5000/\text{QALY}$ can be used to optimally allocate all of the available vaccines. For any initial proportion of susceptibles $\theta_S^0 \geq 0.76$, all 550 vaccines should be used.

Now we consider a more complex scenario of vaccine limitation where vaccines become available in varying quantities over several decision epochs. Let us assume that at time $t = 0$, the policy maker knows with certainty that during the following H decision epochs $\{1, 2, \dots, H\}$, $\{n_1, n_2, \dots, n_H\}$ vaccines will become available. Like before, we assume that any vaccine unused during a period will be lost at the end of the period. Given epidemic state $(X_S(t), X_I(t)) = (x'_S, x'_I)$ and n_t available vaccines at time t , the policy maker must now decide how many susceptibles to vaccinate and also whether to implement the transmission-reducing intervention. To illustrate how these decisions can be made at each decision epoch, let us assume that $H = 2$ (cases with $H > 2$ are solved in a similar fashion).

Table 2. Determining Dynamic Health Policies Using Real-Time Data.

<p>At time $t = 0$</p> <ol style="list-style-type: none"> 1. Obtain estimates for model parameters $(\hat{\mu}, \hat{\tau})$, initial number of susceptibles, \hat{x}_S^0, and a prior distribution for random variable $X_I(1)$. 2. Use Eqs. (7)–(8) to determine the dynamic health policy $\pi^* = [a^*(\cdot)]$. 3. Find the optimal number of susceptibles to immunize, z_0^*, using Eq. (10). 4. Increment time to $t = 1$. 5. Update the number of susceptibles at time $t = 1$: $\hat{x}_S^1 = \hat{x}_S^0 - z_0^* - \hat{x}_I^1$. <p>While $\hat{X}_I(t) > 0$:</p> <ol style="list-style-type: none"> 1. Given the availability of vaccine, use Eq. (8) or Eq. (12) to find the optimal number of susceptibles to immunize, z_t^*, and the optimal transmission-reducing intervention, a_t^*. 2. At the end of period t, use the real-time data $\{\hat{I}_1, \hat{I}_2, \dots, \hat{I}_t\}$ to update the model parameters (using for instance, maximum likelihood estimation) and the dynamic health policy $\pi^* = [a^*(\cdot)]$, using Eqs. (7)–(8). 3. Increment time: $t \leftarrow t + 1$. 4. Use Table 1 to update the epidemic state at time t.

doi:10.1371/journal.pone.0024043.t002

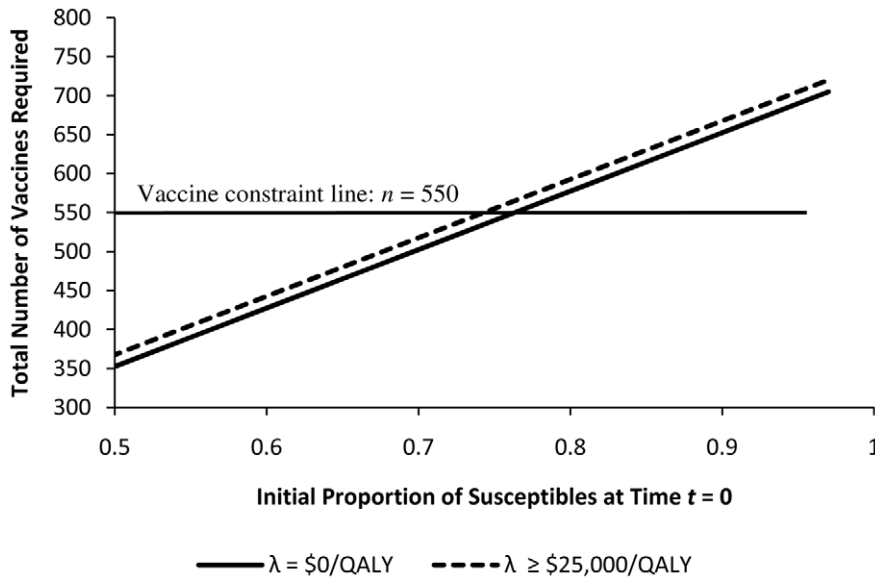


Figure 3. Number of Vaccines Required. Each line specifies the optimal number of vaccines required by different WTP for health. For a given initial number of susceptibles, the policy maker can use this figure to find the WTP for health whose corresponding policy satisfies the vaccine constraint. doi:10.1371/journal.pone.0024043.g003

The optimal allocation is determined iteratively as follows. We start by finding the optimal decisions for the final decision epoch for which vaccines are supplied (i.e. $t = H = 2$ in this example). If at decision epoch $t = 2$, the epidemic is at state $(X_S(2), X_I(2)) = (x_S^2, x_I^2)$ with n_2 vaccines available, the optimal recommendations for vaccination and transmission-reducing intervention are determined by solving the problem:

$$v_2^*(x_S^2, x_I^2) = \max_{\substack{0 \leq z_2 \leq n_2 \\ a_2 \in \{0,1\}}} \hat{v}(x_S^2, x_I^2; a_2, z_2), \quad (13)$$

where $\hat{v}(x_S^2, x_I^2; a_2, z_2)$ is calculated by Eq. 11. Next, we step backward in time and determine the optimal decisions for epoch $t = 1$. If at time $t = 1$, the epidemic is at state $(X_S(1), X_I(1)) = (x_S^1, x_I^1)$ with n_1 vaccines available, the optimal decisions at time $t = 1$ is determined by:

$$v_1^*(x_S^1, x_I^1) = \max_{0 \leq z_1 \leq n_1} r(x_S^1, x_I^1; a_1) - pz_1 + \lambda \sum_{(x_S^2, x_I^2) \in \Omega'(z_1)} P(x_S^2, x_I^2 | x_S^1, x_I^1; a_1) v_2^*(x_S^2, x_I^2), \quad (14)$$

where $\Omega'(z_1) = \{(x_S^2, x_I^2) \in \Omega | x_S^2 = x_S^1 - z_1\}$ and $v_2^*(x_S^2, x_I^2)$ is already calculated in problem (13).

Now to employ this new dynamic health policy that was generated under assumptions of known vaccine constraints $\{n_1, n_2\}$, we take the following steps (note that we assume here that no vaccine is available at time $t = 0$):

Step 1: At time $t = 1$, for the observed state $(\hat{x}_S^1, \hat{x}_I^1)$ and n_1 available vaccines, use problem (14) to find the optimal number of susceptibles to vaccinate and the transmission-reducing intervention to implement.

Step 2: Update the epidemic state using the observed number of new cases during period 1, \hat{I}_1 (see Table 1).

Step 3: At time $t = 2$, for the observed state $(\hat{x}_S^2, \hat{x}_I^2)$ and n_2 available vaccines, use problem (13) to find the optimal number of susceptibles to vaccinate and the transmission-reducing intervention to implement.

Step 4: For the remaining decision epochs $t \geq 3$, use Table 1 to update the epidemic state and then use problem (8) to find the optimal transmission-reducing intervention to employ.

Finally, we discuss how budget constraints can be accounted for while generating dynamic health policies. From a methodological perspective, ensuring that budgetary limitations are not exceeded while attempting to control an epidemic is more challenging primarily due to the stochastic nature of disease spread that leads to high variance for the expected costs incurred. However, if we assume that the policy maker is mainly interested in keeping the *expected* cost incurred during epidemic lower than a constant budget threshold, Figure 4 can be used to select the health policy which satisfies such a constraint.

Suppose that for our population, after vaccination phase at $t = 0$, the proportion of susceptibles at $t = 1$ is reduced to 0.7, and that the policy maker sets the budget threshold at $b = \$50,000$. Then, Figure 4 indicates that for WTP $\lambda = \$100,000/\text{QALY}$ the expected cost incurred during epidemic remains below b at $\$50,000$. Therefore, the policy maker should use the health policy generated by setting the willingness-to-pay $\lambda = \$100,000/\text{QALY}$ to guide decision making during the epidemic.

Evaluating the Effect of Dynamic Health Policies in Controlling the Spread of Influenza

To study the effect of employing the dynamic health policies, we built a simulation model for influenza spread in the population described above. Figure 5 displays the expected number of new infections during each period when $\alpha_T = 40\%$ and policies corresponding to different WTP are employed, given the

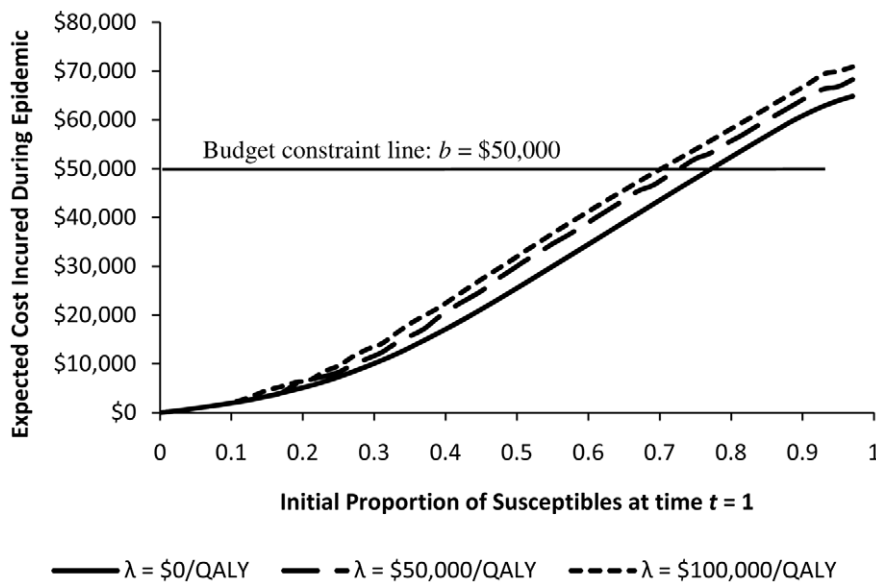


Figure 4. Expected Budget Required During Epidemic. Each curve specifies the expected budget required by different WTP for health. For a given initial number of susceptibles, the policy maker can use this figure to find the WTP for health whose corresponding policy satisfies the budget constraint.
doi:10.1371/journal.pone.0024043.g004

realistic assumption that no effective vaccine was available before the start of epidemic. If vaccines were available in sufficient quantities to employ the and the vaccine recommendations in Figure 5, the epidemic would be averted and there would be no need to employ dynamic policies. The incidence of disease for these policies in each period were statistically different with each other at 1000 simulation runs. As shown in Figure 5, as the willingness-to-pay for health increases, the expected number of individuals infected during the epidemic is reduced since the policies corresponding to the higher willingness-to-pay tend to be more aggressive in implementing the transmission-reducing intervention.

Evaluating the Assumption of the Observability of New Cases

In previous section, we assumed that the policy maker is able to accurately measure the number of new infections occurring in each period. This is a strong assumption that will be violated for many infectious diseases (including influenza) where diagnosis is difficult and for which the number of reported cases is an underestimate for the actual number of infections. To examine the sensitivity of the performance of the generated dynamic health policies to the assumption that all cases are observed, we assume that the policy maker’s observation of the number of cases during

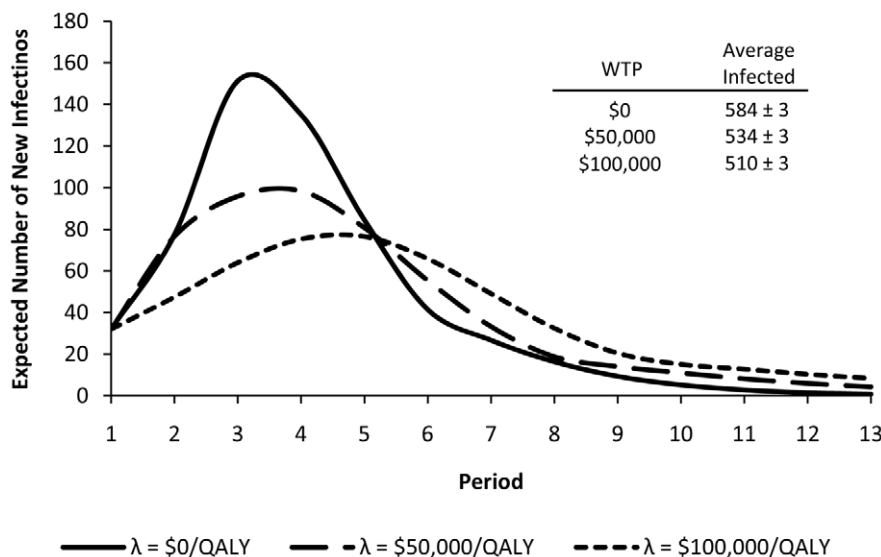


Figure 5. Effect of dynamic health policies on controlling the spread of influenza when no vaccine is available. As the willingness-to-pay for health increases, the expected number of individuals infected during the epidemic is reduced since the policies corresponding to the higher willingness-to-pay tend to be more aggressive in implementing the transmission-reducing intervention.
doi:10.1371/journal.pone.0024043.g005

period $[t, t + \Delta t]$ is only a portion of the true number of cases; that is $\hat{x}_I^{t+1} = (1 + e(t))I(t)$, where $e(t)$ is the percent error in identifying the number of new cases during period $[t, t + \Delta t]$.

If the error term $e(t)$ is a constant in each period, i.e. $e(t) = \eta_t$, and known to the policy maker, the number of susceptibles and the number of infectives in Table 1 can be calculated, respectively, as $\hat{x}_S^t = \hat{x}_S^{t-1} - (1 + \eta_{t-1})\hat{I}_{t-1}$ and $\hat{x}_I^t = (1 + \eta_{t-1})\hat{I}_{t-1}$. Hence, the constant observation error $e(t) = \eta_t$ can be easily corrected in order to preserve the optimality of dynamic health policies under this error setting.

Now let us assume that the error term $e(t)$ has the form $e(t) = \eta_t + \varepsilon$, where η_t is a constant known to the policy maker and ε is a normally distributed noise with mean zero and standard deviation $\sigma(\varepsilon)$. The effect of the constant term η_t can be corrected like before by using equations $\hat{x}_S^t = \hat{x}_S^{t-1} - (1 + \eta_{t-1})\hat{I}_{t-1}$ and $\hat{x}_I^t = (1 + \eta_{t-1})\hat{I}_{t-1}$ to update the current information about the epidemic state in Table 1. The effect of the random noise cannot be corrected; yet, the impact of this noise on the performance of the dynamic health policies can be investigated through simulation. Figure 6 shows the effect of uncertainty around the number of observed cases on the capability of dynamic health policies in controlling the epidemic in our population, for different WTP for health. We find that as the standard deviation of noise, $\sigma(\varepsilon)$, increases, the total number of individuals affected by the epidemic also increases; however, this increase is linear with a low rate which implies that for our population the performance of the generated dynamic health policies is not highly sensitive to the assumption of observability of new cases.

If this simulation analysis shows significant sensitivity of the performance of dynamic health policies to the assumption of observability of new cases, or the policy maker believes that the probability distribution of noise ε varies over time (for instance, if vigilance of testing diagnosis increases or decreases over time),

then more advanced optimization tools can be employed. This issue is briefly discussed in the Discussion section.

Discussion

The emergence of novel human pathogens (e.g. H1N1 and H5N1 influenza, SARS) and their devastating health and financial consequences on affected populations have highlighted the need for developing methods which allow real-time selection of health interventions to control the epidemic while effective vaccines are not available or not present in sufficient quantities to prevent disease spread. We refer to policies informed by such methods as *dynamic health policies* which are intended to allow for *real-time* recommendations to be made in response to changing disease and population characteristics as well as the availability of resources.

In contrast to most existing approaches for identifying optimal strategies for infectious disease control which use simulation or mathematical models of disease spread to compare the performance of a *limited* number of *pre-determined* health policies, we proposed the use of “dynamic programming” [23] to characterize and identify optimal dynamic health policies. We demonstrated how a Markov decision process [16] can be employed to find optimal dynamic health policies for a simple model of influenza epidemic, in which two types of interventions may be available during the epidemic to control the influenza spread: (1) vaccination, and (2) a transmission-reducing intervention, such as social distancing. The generated dynamic health policies help the policy maker to determine (1) how to allocate vaccines when they become available, and (2) whether the transmission-reducing intervention, such as school closure, should be employed or lifted given the number of susceptibles and infectives at any point of time.

While we used a discrete-time Markov decision process, a number of other methodologies have also been proposed and

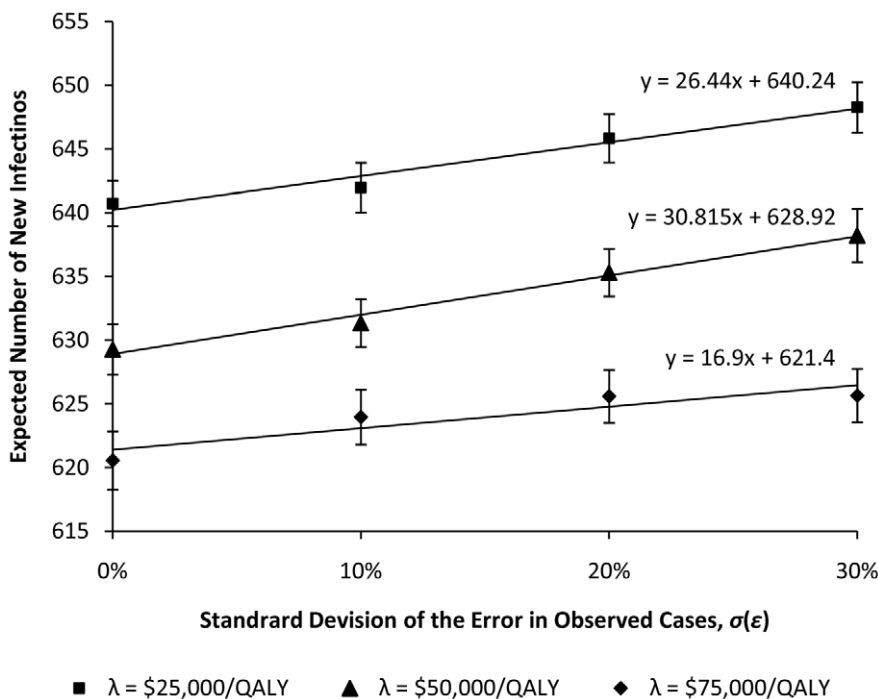


Figure 6. Expected total number of new infections versus the standard deviation of the error in observing the number of new cases. As the standard deviation of noise increases, reflecting imperfect surveillance capacity, the total number of individuals affected by the epidemic also increases. However, this increase occurs linearly at a low rate, which implies that for our population the performance of the generated dynamic health policies is not extremely sensitive to the assumption of observability of new cases. doi:10.1371/journal.pone.0024043.g006

developed to determine or approximate optimal dynamic health policies for controlling emerging epidemic. For instance, Lefevre [24] used a continuous-time Markov decision model, Merl et al. [2] developed a statistical framework and Ludkovski and Niemi [21] developed a simulation-based model for dynamic determination of optimal policies for emerging epidemics. Undoubtedly, comparing the effectiveness of these methodologies when employed in real practice merits a separate research study.

The influenza model proposed in this paper is very simple and not intended to realistically model disease spread or be used directly to guide the selection of interventions. We present it only for illustration of our proposed approach for dynamic decision making. The model makes several simplifying assumptions that were required for using an MDP to generate the optimal health policies: first, we require that the number of new infections during each period is observable by the policy maker, and second, we assume that an infectious individual interacts with the rest of the population only during the next period and then is effectively removed (treated or isolated) from the population. Relaxing these two assumptions will mean the state of the epidemic is unobservable; yet, a probability *belief* about the epidemic state can often be generated as new data become available. If we relax these assumptions, we can use generalized discrete-time Markov models for infectious diseases proposed in [19] and use partially observable Markov decision process (POMDP) [25] to characterize optimal health policies. However, as the number of states required to model disease spread increases, MDP and POMDP rapidly lose their efficiency. In these cases, approximate dynamic programming [22] using a simulation model may be employed to identify optimal dynamic health policies. Accordingly, the framework we propose here can potentially be extended to inform decision-making for control of a pathogen with a more complex natural history such as tuberculosis or to design interventions that

consider distinct responses targeted toward different risk groups. Such extensions are attractive topics for future research.

As a final note, for the successful implement of the dynamic health policies in practice, the mathematical or simulation model of the disease spread along with the optimization technique used for finding the dynamic health policies must be coupled with an surveillance system that can supply data to estimate the parameters of the underlying model and to provide knowledge on the state of the epidemic. Although dynamic optimization techniques are capable of handling noisy observations, inaccuracies in the surveillance and reporting system may result in suboptimal policies, further underscoring the tremendous importance of public health surveillance in defining responses to epidemics.

Supporting Information

Text S1 Calculating the Transition Probabilities for the Markov Model of Influenza Spread. Approximating the State Space for Influenza Spread. Additional Information for the Illustrative Example. (PDF)

Acknowledgments

The authors would like to thank Marc Lipsitch for his comments and suggestions on this work.

Author Contributions

Conceived and designed the experiments: RY TC. Performed the experiments: RY TC. Analyzed the data: RY TC. Contributed reagents/materials/analysis tools: RY TC. Wrote the paper: RY TC.

References

- Wallinga J, van Boven M, Lipsitch M (2010) Optimizing infectious disease interventions during an emerging epidemic. *Proceedings of the National Academy of Sciences* 107: 923–928.
- Merl D, Johnson L, Gramacy R, Mangel M (2009) A statistical framework for the adaptive management of epidemiological interventions. *PLoS One* 4: e5087.
- Dushoff J, Plotkin J, Viboud C, Simonsen L, Miller M, et al. (2007) Vaccinating to protect a vulnerable subpopulation. *PLoS Med* 4: e174.
- Goldstein E, Apolloni A, Lewis B, Miller J, Macauley M, et al. (2010) Distribution of vaccine/antivirals and the ‘least spread line’ in a stratified population. *Journal of the Royal Society Interface* 7: 755–764.
- Patel R, Longini I, Halloran ME (2005) Finding optimal vaccination strategies for pandemic influenza using genetic algorithms. *Journal of Theoretical Biology* 234: 201–212.
- Tanner M, Sattenspiel L, Ntamo L (2008) Finding optimal vaccination strategies under parameter uncertainty using stochastic programming. *Mathematical Biosciences* 215: 144–151.
- Arinaminpathy N, McLean A (2008) Antiviral treatment for the control of pandemic influenza: some logistical constraints. *Journal of the Royal Society Interface* 5: 545–553.
- Dimitrov N, Goll S, Meyers L, Pourbohloul B, Hupert N (2009) Optimizing tactics for use of the US antiviral strategic national stockpile for pandemic (H1N1) Influenza, 2009. *PLoS Curr Influenza*.
- Ferguson N, Cummings D, Fraser C, Cajka J, Cooley P, et al. (2006) Strategies for mitigating an influenza pandemic. *Nature* 442: 448–452.
- Ferguson N, Cummings D, Cauchemez S, Fraser C, Riley S, et al. (2005) Strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature* 437: 209–214.
- Halloran M, Ferguson N, Eubank S, Longini I, Cummings D, et al. (2008) Modeling targeted layered containment of an influenza pandemic in the United States. *Proceedings of the National Academy of Sciences* 105: 4639–4644.
- Flahault A, Vergu E, Coudeville L, Grais R (2006) Strategies for containing a global influenza pandemic. *Vaccine* 24: 6751–6755.
- Longini I, Jr., Nizam A, Xu S, Ungchusak K, Hanshaoworakul W, et al. (2005) Containing pandemic influenza at the source. *Science* 309: 1083–1087.
- Fedson D (2003) Pandemic influenza and the global vaccine supply. *Clinical Infectious Diseases* 36: 1552–1561.
- Webby R, Webster R (2003) Are we ready for pandemic influenza? *Science* 302: 1519–1522.
- Puterman ML (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York: John Wiley & Sons, Inc.
- Grassly NC, Fraser C (2008) Mathematical models of infectious disease transmission. *Nature Reviews Microbiology* 6: 477–487.
- Stimmet AA, Mullahy J (1998) *Net Health Benefits: A New Framework for the Analysis of Uncertainty in Cost-Effectiveness Analysis*. *Medical Decision Making* 18: S68–S80.
- Yaesoubi R, Cohen T (In Press) Generalized Markov models of infectious disease spread: A novel framework for developing dynamic health policies. *European Journal of Operational Research*.
- Anonymous (1978) Influenza in a boarding school. *British Medical Journal*. 587 p.
- Ludkovski M, Niemi J (2010) Optimal dynamic policies for influenza management. *Statistical Communications in Infectious Diseases* 2.
- Powell W (2007) *Approximate Dynamic Programming: Solving the curses of dimensionality*. Wiley-Interscience.
- Winston WL (2003) *Operations Research: Applications and Algorithms*. Duxbury Press, fourth edition.
- Lefevre C (1981) Optimal control of a birth and death epidemic process. *Operations Research* 29: 971–982.
- Sondik E (1978) The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs. *Operations Research* 26: 282–304.