



Service Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Multiregional Dynamic Vaccine Allocation During an Influenza Epidemic

Anna Teytelman, Richard C. Larson

To cite this article:

Anna Teytelman, Richard C. Larson (2013) Multiregional Dynamic Vaccine Allocation During an Influenza Epidemic. Service Science 5(3):197-215. <https://doi.org/10.1287/serv.2013.0046>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2013, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Multiregional Dynamic Vaccine Allocation During an Influenza Epidemic

Anna Teytelman, Richard C. Larson

Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
{teytanna@alum.mit.edu, rclarson@mit.edu}

We consider a sequential decision problem in national health policy: the weekly deployments of limited influenza vaccine doses, as they become available, to the various geographical regions of the United States during a pandemic event. As with the 2009 H1N1 pandemic flu, we assume that the progression of flu infection varies from region to region, with some regions starting their flu waves weeks before others. We also assume that vaccine doses only become available after flu waves have already started in some regions. Whereas the traditional deployment of vaccines is in direct proportion to resident population, without regard to flu status in any region, we show that a new policy that dynamically considers the status of the various regional flu waves can dramatically reduce the incidence of flu infection over the entire country. The method uses mathematical models of flu spread and requires capturing real-time data on flu incidence.

Key words: influenza modeling; H1N1; vaccine allocation; pandemic response

History: Received December 1, 2012; Received in revised form March 18, 2013; Accepted March 25, 2013 by Paul Maglio. Published online in *Articles in Advance* June 12, 2013.

Introduction

Influenza vaccines are highly effective in providing personal immunity from a particular virus (Hardelid et al. 2011, Harris et al. 2010). They have been used to battle the devastating effects of pandemic influenza since the middle of the 20th century. However, because of delays in vaccine creation, manufacture, and distribution, careful thought must be used in deploying vaccine doses once they start to become available. In the United States, vaccines routinely take six months to manufacture once a novel flu virus has been identified, and only then can shipments in bulk quantities commence (Centers for Disease Control and Prevention 2012). As a result, vaccines may only begin to become available when an outbreak is in full force or even after most of the first wave has subsided. The first vaccine shipments are usually much smaller than vaccine demand. Limited supplies must be allocated carefully to ensure that they are being used with the highest possible beneficial effect. As an additional complication, vaccines often carry perceived risks (Hardelid et al. 2011), so full population compliance is not guaranteed even when vaccines are available—even the compliance of U.S. healthcare workers has been less than 50% since 1989 (Centers for Disease Control and Prevention 2010b). Given many individuals' reluctance to accept the flu vaccine, we are fortunate that theoretical research on herd immunity suggests that the percentage of people needed to be vaccinated to provide adequate protection for the entire population is only somewhere between 30% and 50% (Hill and Longini 2003, Teytelman 2012).

The effective and timely use of limited incoming supplies of vaccines is of utmost importance and is the focus of this paper. We discuss the problems of vaccine allocation in the United States and propose methods (“algorithms”) for real-time vaccine allocation during a major event such as the H1N1 influenza pandemic of 2009–2010.

The 2009–2010 H1N1 pandemic was the first worldwide influenza event for which global and local authorities had the tracking resources available to gather extensive data on disease spread and the various interventions used to mitigate the pandemic's effect. The United States has a single vaccine-dispensing body, the Centers for Disease Control and Prevention (CDC), and only this one body need accept the suggested new algorithms for vaccine distribution. This “single decision maker, single dispenser” structure makes feasible what we are proposing here.

The World Health Organization (WHO) and CDC identified the novel H1N1 virus in Mexico and the United States in early April 2009. However, the virus did not spread rapidly in the United States until schools started opening in late summer. Because of their early school openings in August (Centers for Disease Control and Prevention 2009), Georgia, Alabama, and other southeastern states experienced the major brunt early in the U.S. outbreak. By the time large quantities of vaccines had been manufactured and shipped by the CDC in October (Centers for Disease Control and Prevention 2009), most southeastern states had already experienced a peak followed by a decline in the prevalence of influenza. At that time, the northern states were still in the early

stages of their outbreaks and were ready to vaccinate a significant portion of their population prior to the peak of the epidemic. Despite these considerable regional differences in the progressions of flu waves, the CDC initially delivered vaccines to states strictly on a per-capita basis (Centers for Disease Control and Prevention 2010a), with each state receiving vaccine in direct proportion to its population in what was determined to be the fairest distribution scheme.

Benefits derived from the shipped vaccines varied considerably from state to state, as discussed by Finkelstein et al. (2011) and Teytelman and Larson (2012). States that received vaccines early with respect to their epidemic peak had a higher uptake rate in the population. Residents were more willing to accept the vaccine when the perceived risk of the epidemic was still high, and vaccines that were administered early were more effective in preventing the spread of the outbreak and limiting the total number of infections. In October 2009, some southeastern states—where flu waves were nearly over—were using as little as 25% of the vaccine allocated to them in the first few weeks of the shipment process and did so with little beneficial effect. Other states were administering vaccines with higher beneficial effect, but they were limited by the constrained supply (Larson and Teytelman 2012).

The dynamic decision problem we face is this: *In the days and weeks following the first limited shipments of vaccine, as more and more vaccine doses become available, and operating in a large region such as the United States where the progression of the flu wave varies markedly over geographical regions, how does one allocate vaccine over time to the various regions in the most beneficial manner?*

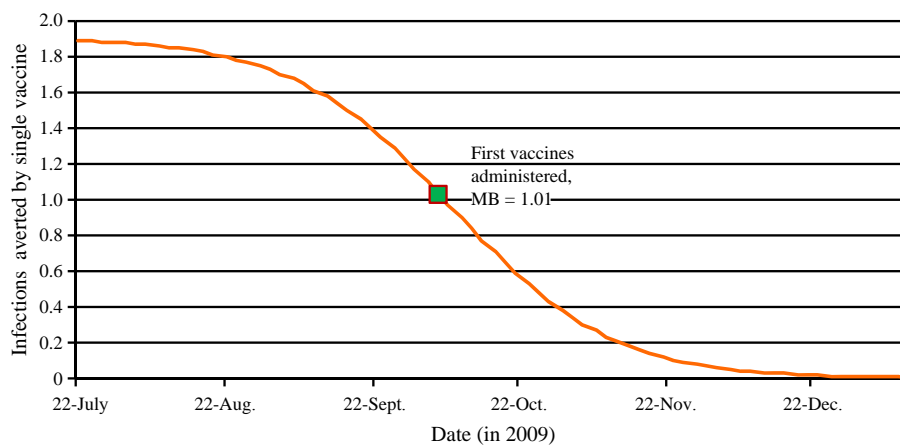
For the single-region problem, where geographic flu wave variability is not a major issue, considerable research has been performed evaluating different vaccine-targeting strategies. Suggestions range from prioritizing those persons at risk (Chowell et al. 2009, Longini and Halloran 2005, Patel et al. 2005) to vaccinating those individuals that contribute most to virus spread (Longini et al. 2005, Nigmatulina and Larson 2009). Although these nonspatial issues are important to consider, we focus on a new higher-level algorithm to allocate limited supplies of vaccines to various spatially disparate regions based on information about the progression of the flu wave in each of the respective regions. Then, once a given amount of vaccine is delivered to a region as a result of our proposed methods, previous person-specific methods for deploying that vaccine in the region can and should be used.

Motivation

The modeling approach of Larson and Teytelman (2012) provides the motivation for our methodology. Using the model described in that paper, we display in Figure 1 the marginal benefit of a hypothetical single vaccine dose if it were administered at different points in 2009 in the state of Oklahoma.

In this example, if Oklahoma receives vaccine doses early in the pandemic, a single dose will, on average, avert flu infection in almost two people. This average includes the person being vaccinated as well as the expected number of people who would contract the disease directly or indirectly from the vaccinated person had she not been vaccinated. On the other hand, if a region receives vaccine doses late, they are virtually wasted. Note the inverted S shape of the curve in Figure 1. Although the curve is monotonically decreasing, it is relatively flat in July and early August as well as in December. This means that the relative effectiveness

Figure 1. Marginal Benefit (MB) of a Single Vaccine Dose Administered at Different Times in Oklahoma Based on Data from the 2009–2010 H1N1 Pandemic



of a vaccination would be virtually unchanged if we were to administer a dose in July or the beginning of August. However, a week's delay in September would decrease the effectiveness of a vaccine significantly. The flatness in December indicates the virtual ineffectiveness of the vaccine if administered that late. We thus arrive at what we label a *critical period*, ranging from the middle of August until the end of November. It is intuitive to observe that the peak of the epidemic curve actually occurred on the week of October 10th in Oklahoma, right in the middle of the critical period (Larson and Teytelman 2012).

In a multiregion problem, where different regions are at different stages of the epidemic, we want to focus on these critical time periods, which are usually different for each region. In general, it is best to administer the vaccines no later than early in the critical period. We want the vaccines administered before the precipitous fall in effectiveness during the critical period. In trading off various deployments, states that have not yet entered their respective critical periods will not be too disadvantaged by delaying their vaccinations.

Basic Allocation Problem

If our proposed methods are to be implemented, the multiregion vaccine allocation problem has to be addressed by the CDC several times throughout the outbreak; we suggest once each week for the duration of the outbreak. At each decision point, the CDC needs to make decisions in two steps.

Step 1: Set Up and Continuously Update Flu Wave Curves. For each region, the CDC needs to develop and update approximate and partial flu wave curves. Each partial curve will display the history to date (i.e., the number of influenza-like illnesses reported, or ILIs) and make a projection into the future of the flu wave curve *in the absence of vaccines*. These data will likely contain errors of under- and overreporting of illness, and they are often delayed. Underreporting is due to infected individuals not presenting themselves to physicians. Overreporting is often due to the “worried well” presenting themselves or their family members to physicians. However, data inaccuracies should not preclude the construction and concomitant parameter estimation required to initiate region-specific models that will become more accurate as additional data arrive in the following weeks. Parameter estimation will also become more accurate as more data become available. Early estimation of the basic reproductive number R_0 is often difficult because of small sample sizes and noisy data, so the use of a best-practice global estimate of R_0 should suffice until more local data become available to estimate a region-specific value for R_0 (Fraser et al. 2009). Although real-time data are often uncertain and difficult to manage, new technologies are making healthcare surveillance and modeling realistic tools for future outbreaks. An intriguing example is Google, now able to track influenza progression using Internet searches (Carneiro and Mylonakis 2009, Harder et al. 2011).

Step 2: Deploy Vaccine Doses to Regions. In this second stage, we assume that we have a calibrated and updated flu wave model for each region. We must now make decisions to allocate vaccines to obtain the greatest beneficial effect for the entire country. We will propose algorithms that suggest the numbers of available vaccine doses to be shipped to each region.

After the vaccines arrive in a given region, the state takes it upon itself to organize the logistics of locally distributing the vaccines. Once the vaccines are administered, there is a delay of one to two weeks until the vaccine provides full immunity.

Perspective: “Equity vs. Effectiveness”

As discussed above, the current CDC method of deploying vaccines during a flu epidemic is on a per-capita basis, i.e., in direct proportion to each region's population, independent of the status of the flu waves in the various regions. The implied objective function is to “maximize equity”—that is, to give regions the amount of vaccine needed to cover an equal percentage of their respective populations.

Our approach focuses on health outcomes and considers the United States in its entirety. Our objective function is to deploy vaccine doses in order to minimize the expected number of influenza infections over the entire country.

One may dramatically visualize the two competing objectives using an old-fashioned fire bucket brigade, set up using buckets of water to fight ongoing fires in a number of contiguous neighborhood houses. In our approach, we seek to apply the delivered water to burning houses so as to minimize the expected fire damage to the entire neighborhood. Each bucket of water when delivered to the zone of the fires is applied to that fire or fires so as to have the most beneficial effect in terms of minimizing future fire damage. In the “equity” approach, water is deployed according to house size; that is, buckets of water are administered to houses in direct proportion to house size alone, independent of the status of the fire in each house—even if the fire has already been extinguished.

Problem Formulation

To formulate this problem precisely, we will use a discrete-time influenza spread model defined by Teytelman and Larson (2012), where we call the unit of time a “day.” However, the allocation heuristics presented here may just as well be adapted to any influenza spread model.

We say that we have R regions, each having a population n_i . We also have a known vaccine schedule vector \vec{v} with length K , where K is the total number of our modeled days. The scalar $v(k)$ is total number of available vaccine doses on day k . A deciding body (e.g., the CDC) must make an allocation decision on each of T decision points for which $v(k) > 0$.

A decision consists of R vectors \vec{v}_j for $j = 1, \dots, R$, such that $\sum_{j=1}^R \vec{v}_j = \vec{v}$. Each vector \vec{v}_j represents the number of vaccine doses allocated to region j , and $v_j(k)$ is the number of vaccine doses allocated to region j at time k . Let \vec{e}_k be the vector such that

$$e_k(m) = \begin{cases} 0 & \text{if } m \neq k, \\ 1 & \text{if } m = k. \end{cases}$$

To include history, let \vec{v}_i^* be the vector of all vaccines already shipped prior to decision day t . Also, let $E_i(t)$ be the number of infections on day t in region i , let $I(t)$ be the observed value of the epidemic curve on day t , and let $I(\vec{v}_i)$ be the expected total number of infections occurring in region i if it received the vaccines specified by \vec{v}_i .

Our goal is to make a decision $v_j(t)$ for $j = 1, \dots, R$ on each day t for which $v(k) > 0$ in such a way that the expected sum of total infections over all regions, $\sum_{j=1}^R I(\vec{v}_j)$, is minimized. In cases of particularly deadly epidemics, similar analyses may be performed with different objective functions, such as minimizing the total number of deaths or the total economic impact of an epidemic.

Now that we have formalized the quantities needed for our decision-making process, we proceed with suggesting allocation heuristics to evaluate. Recall that a heuristic is an approximate, usually intuitive, way to try to solve a problem—it is not “mathematically optimal” in a rigorous way.

Proposed Heuristics

Pro Rata Heuristic

The pro rata heuristic is the one currently used by the CDC to allocate vaccine doses in direct proportion to populations. Specifically,

$$v_i(t) = v(t) \cdot \frac{n_i}{\sum_{i=1}^R n_j}$$

Prepeak Heuristic

A simple, “quick and dirty” way to increase the effectiveness of vaccines is to allocate all vaccines only to those states that have not yet reached the peak of infection. That is, to allocate to states for which the epidemic curve is currently increasing, vaccines should be distributed proportionally within that set of states:

$$v_i(t) = \begin{cases} v(t) \cdot \frac{n_i}{\sum_{j \in A} n_j} & \text{for } i \in A, \\ 0 & \text{otherwise,} \end{cases}$$

where $i \in A$ iff $E_i(t) > E_i(t-1)$.

This all-or-nothing approach is problematic. States just past their epidemic curve peaks may very well benefit from some amounts of vaccine. Moreover, slight variations in the observed epidemic curves as a result of the probabilistic nature of an epidemic could mischaracterize states that have not yet truly reached the peak of the epidemic but whose flu wave curves incorrectly suggest that they have; such states would not be allocated vaccines during the growth stages of the outbreak when vaccines are most effective. Such mistakes can be very costly and are also bound to elicit strong disagreement from the states passed over for vaccine shipments.

Greedy Heuristic

The most intuitively straightforward approach to allocating vaccines is to allocate the vaccine doses one by one to the region with the highest marginal benefit for the next vaccine. To use the greedy heuristic on day t , we ignore all information about any future vaccines and make a decision based only on vaccines already allocated prior to day t and those available today. For each vaccine, we predict the marginal benefit of allocating this

vaccine to all regions $1, \dots, R$ using some influenza spread model. We then allocate this vaccine to the region that would obtain the highest marginal benefit from the dose.

Although this greedy algorithm performs better than the pro rata heuristic, the greedy property causes suboptimal solutions. Consider a two-region problem, where Region 1 is experiencing the peak of its epidemic early on—say, on day 5—and the other is set to experience the outbreak much later, on day 25. Suppose also that only two shipments of vaccines are available, one on day 5 and the other on day 20. Intuition tells us that we should give most of the first shipment to the first region and all of the second shipment to the second region. However, on day 5, the marginal benefit of most vaccines for the second region will be slightly higher than the marginal benefit for the first region, so the greedy algorithm will allocate the majority of the first shipment to the second region as well. The greedy algorithm does not take into consideration later interventions and thus is likely to make ineffective decisions early on.

Critical Period Heuristic

To illustrate the shortcoming of the greedy heuristic, consider Figure 2, where we superimpose the marginal benefit curves of two regions. The shapes of the two curves are identical, but Region 1 experiences the outbreak three weeks prior to Region 2. Consider a time near the beginning of September, at the start of the critical period for Region 1. Whereas the marginal benefit of a single vaccine dose is slightly higher in Region 2, the time derivative or slope of the Region 2 curve is also relatively low compared with that of Region 1. Region 2 can afford to wait while we focus on Region 1 during its highly time-sensitive period.

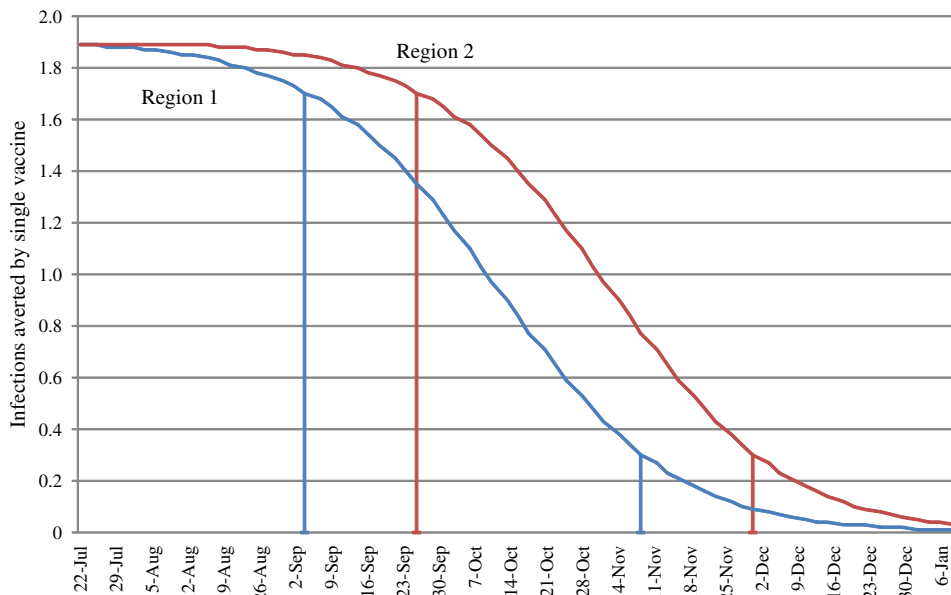
Note that the marginal benefit of a single vaccine in region i at time t is $MB_i(t) = I(\vec{v}_i + \vec{e}_i) - I(\vec{v}_i)$. We then say that a decision day t is in the critical period with threshold ϵ , $CP_{i,\epsilon}$, if $-(dMB_i(t)/dt) > \epsilon$. Thus, for this heuristic we set some ϵ and allocate vaccines on day t as

$$v_i(t) = \begin{cases} v(t) \cdot \frac{n_i}{\sum_{j \in B} n_j} & \text{for } i \in B, \\ 0 & \text{otherwise,} \end{cases}$$

where $i \in B$ iff $t \in CP_{i,\epsilon}$.

Although this heuristic takes into account the future information available at time t , it is still somewhat simplistic in that it also follows the all-or-nothing approach. Some states are given significant preference over others even though the difference in the effect of a vaccine in those states might be fairly close. This is both hurtful to our objective function of minimizing the total number of infections and infeasible in a real-life situation because it introduces unreasonable inequality between regions. Instead, we propose an iterative algorithm that combines the benefits of both the critical period heuristic and the greedy heuristic by looking at the marginal benefit that can be attained at all possible decision points, now and in the future.

Figure 2. Critical Periods for Two Regions with Epidemics Starting at Different Times



Note. The epidemic for Region 1 starts earlier than the epidemic in Region 2.

Telescope-to-the-Future Switching Algorithm Heuristic

To implement the future-oriented switching heuristic algorithm on day t , we need to make some assumptions about the information that is available to decision makers at each decision point. We assume that the following items are known or estimated on each day t :

- R_0 for all regions;
- $E_i(k)$ for $i = 1, \dots, R$, $k \leq t$;
- the function $I(\vec{v}_i)$ for $i = 1, \dots, R$;
- the already allocated vaccine vector \vec{v}_i^* for $i = 1, \dots, R$; and
- the forecasted vaccine schedule \vec{v} .

Given this information, we make a decision for day t and estimate forecasted decisions for all subsequent decision days—that is, we “telescope” into the future. We calculate $\vec{v}_i(t)$ and $\vec{v}_i(k)$ for $k > t$. The algorithm follows the following steps:

Step 1. We start with a possible allocation (satisfying $v_i(t) > 0 \forall i, t$ and $\sum_{i=1}^R \vec{v}_i = \vec{v}$) for all future decision points.

Step 2. Then, for every decision point $k \geq t$, we calculate the benefit of switching a batch of vaccine doses between every pair of regions. Here, “benefit” is the number of flu infections averted.

Step 3. We then select the best switch over *all possible decision points* and update the allocation accordingly. Note that this may happen in the future and not reflect any actual action on day t by switching a hypothetical batch of vaccine doses on some future date from region i to region j .

Step 4. We repeat Steps 2 and 3 until no more improvements in the objective function can be made.

By continually taking into account future decisions, we obtain a solution that maintains a balance between the myopic and more global heuristics. Our forecast into the future is bound to be flawed, however, so we must repeat this procedure at each day $k > t$, continually updating the forecasted decisions for future time periods. This forecasting technique addresses the possible stability concerns with myopic heuristics. Because current decisions will affect the shape of the epidemic curve in the future, we must include in our model the total number of infections under the projected vaccine allocation over all decision points. That way, we make sure that we take into account the effect of earlier decisions on influenza progression in the future.

Despite the fact that this algorithm terminates with some vaccine allocation for all future decision days, the resulting allocation should only be used for the day on which the algorithm was run. In the long run, we must take into account the real-life changes to the curve that occur as the outbreak progresses. Because our data are incomplete and events outside the scope of vaccine allocation might drastically change the progression of the outbreak, this algorithm must be rerun every time new vaccine shipments become available. On each decision day, the model should use updated parameters fitted to the new data that have come in since the last shipment time the decision was determined.

For the formalization of all the algorithms described in this paper, we direct the reader to a more detailed work describing these algorithms (Teytelman 2012). The optimized algorithm is a modified gradient descent algorithm that converges to a near-optimal solution quickly and efficiently; it is provided in Appendix A.

Testing Strategy

Influenza progression in a community is made difficult by the fact that the observed empirical epidemic curve rarely follows any given theoretical epidemic curve. In fact, day-to-day progression is not only volatile but suffers from inaccuracies of available information. Patient reporting to doctors and hospitals is highly dependent on the public estimation of the risk associated with the current outbreak. To test the effectiveness of various allocation heuristics, we must incorporate such uncertainties. We test the heuristics using Monte Carlo simulation to approximate the expected total number of infections that will occur with the CDC using a given heuristic. For each execution of the simulation, we use a discrete-time model, where for each day we perform the following four steps.

Step 1. We first determine the starting time and the fundamental parameters of epidemic spread in each region. To put our results in context, we used data from the 2009–2010 H1N1 epidemic to infer these parameters. We used the H1N1 data provided by the 50 states to the CDC to fit value for R_0 for each region. We then used these estimated parameters as inputs to a stochastic influenza spread model to generate the next day of outbreak. See Teytelman (2012) as well as Appendices B and C for details of the stochastic implementation of the model and the summary parameters used. The estimated parameters we used included the benefit of hindsight and two years of data gathering and analysis. Naturally, this information was not available as the outbreak was happening, so in this case, we assume that these parameters are hidden from decision makers, who experience

only the epidemic curve. For our current testing purposes, we assume that the epidemic curve information is observable to the CDC as it happens. However, more accurate modeling needs to take into account the delay between the information available and real events.

In an ideal world, all vaccine doses shipped to a region would be immediately administered to patients. In reality, vaccine administration is not perfect. In the 2009 epidemic, the time from vaccine arrival to patient administration varied from state to state, sometimes lasting a few weeks (Hopkins 2011) because of logistical difficulties. In addition, the vaccine acceptance percentage was significantly lower than 100%, mostly because the public’s demand for vaccines waned at the end of the outbreak within each state. For the 11 states analyzed in Larson and Teytelman (2012), the percentage of all delivered vaccine doses actually used by each state ranged from 20% to 40%. This number is lower than what we would expect, but it is not discouraging if we understand that most of the unused vaccine was shipped at the end of the vaccination program, when most of the demand for vaccines has gone down. In the first few weeks of shipments, the administered vaccine percentage was as high as 60% in some states. For our testing framework, we estimated the administered vaccine percentage for each region from the H1N1 data and used this percentage to model a similar regional and public reaction to the vaccine available. Appendix C contains the values of all parameters used in the testing model.

Once the next day is generated, if new vaccine becomes available to be shipped, we move on to testing the decision-making process.

Step 2. The decision makers must now estimate the model parameters to be used for the allocation algorithm using only the data that would be available in a real-time situation.

Step 3. Armed with approximated parameters, the decision makers can now use one of the proposed heuristics to output a set of values $v_i(t)$ that represent the amount of available vaccine that is to be distributed to each region.

Step 4. The decision makers apply their decisions to each region by shipping the vaccines according to the decisions in Step 3. This allocation will now be used in Step 1 when the next generation day is calculated.

Naturally, if there is no vaccine available—that is, $v(t) = 0$ —Steps 2–4 need not be performed, and we simply generate the next day in Step 1.

Testing Results

Recreating the 2009–2010 H1N1 Epidemic

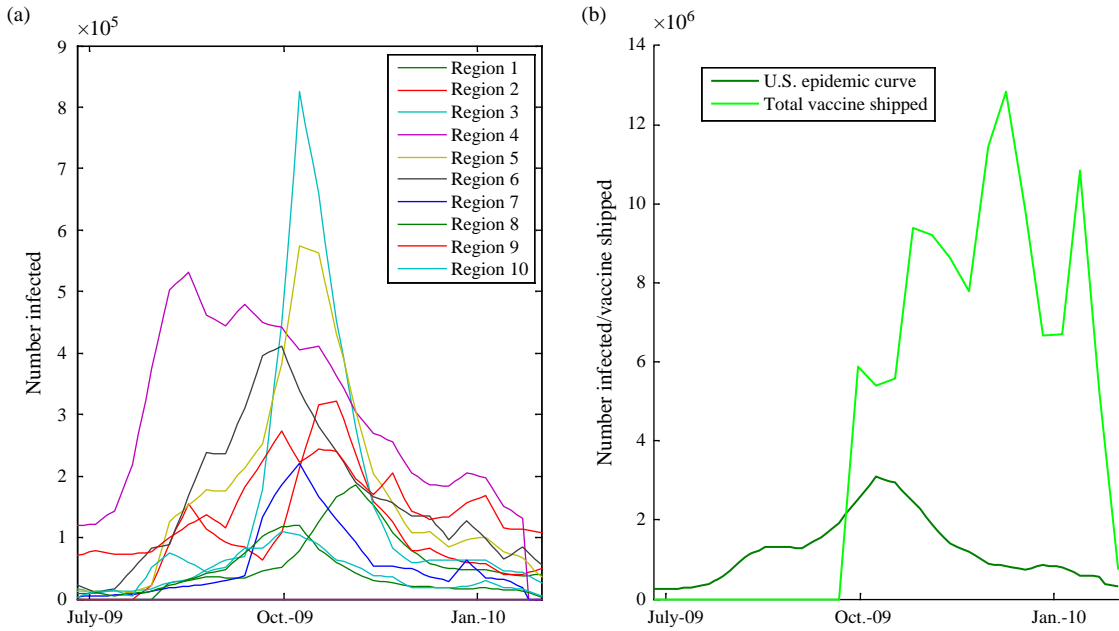
We used a Monte Carlo simulation model; the parameters were fit to the 2009–2010 H1N1 epidemic data as well as other characteristics of the 10 regions of the United States as classified by the CDC (see Table 1).

As in 2009, most of the regions started experiencing signs of a major epidemic in the summer, whereas the vaccines started arriving about three months later—on October 10th. From that point on, decisions to ship the available vaccine doses were made on a weekly basis. Figure 3(a) shows the real estimated epidemic curves for the 10 regions as the outbreak occurred. The epidemic curves were obtained from the percent ILI numbers reported by sentinel sites in the same manner as Larson and Teytelman (2012). We also include the vaccines as they were made available to be shipped by the CDC, as depicted in Figure 3(b).

Table 1. Regions of the United States Used by the CDC for Vaccine Allocation

Region	States
1. New England	Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, Vermont
2. New York Area	New Jersey, New York
3. Mid-Atlantic	Delaware, Maryland, Pennsylvania, Virginia, West Virginia
4. Southeast	Alabama, Florida, Georgia, Kentucky, Mississippi, North Carolina, South Carolina, Tennessee
5. Great Lakes	Illinois, Indiana, Michigan, Minnesota, Ohio, Wisconsin
6. Southwest	Arkansas, Louisiana, New Mexico, Oklahoma, Texas
7. The Plains	Iowa, Kansas, Missouri, Nebraska
8. Rocky Mountains	Colorado, Montana, North Dakota, South Dakota, Utah, Wyoming
9. West	Arizona, California, Hawaii, Nevada
10. Pacific Northwest	Alaska, Idaho, Oregon, Washington

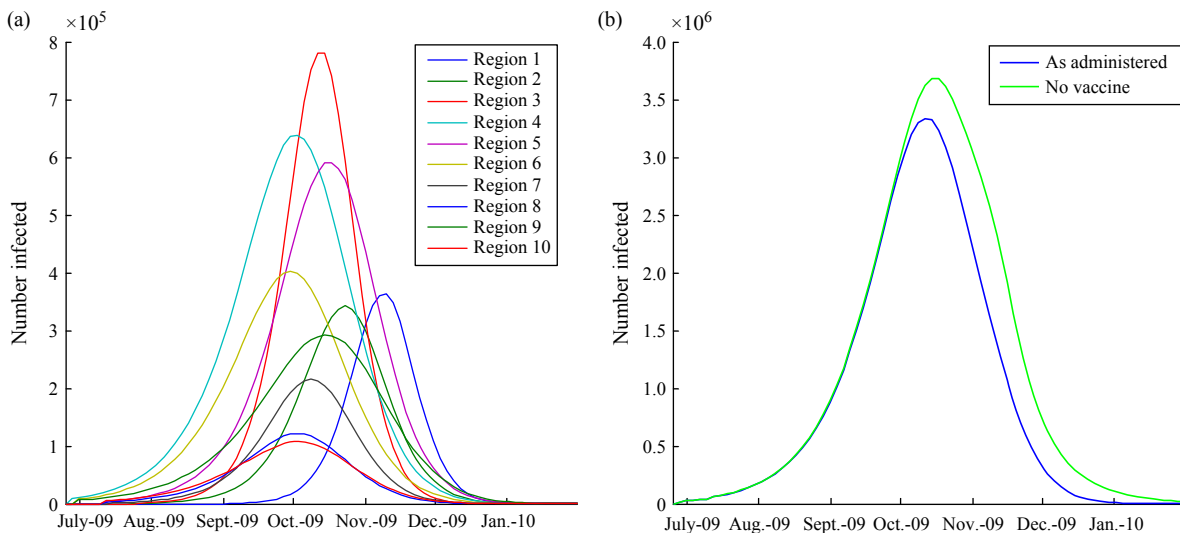
Figure 3. Estimated Epidemic Curves by Region (a) and Vaccine Doses Shipped and Total Epidemic Curves for the United States (b) Based on 2009–2010 H1N1 Epidemic Data



Unfortunately, vaccines started being shipped just before the peak of the infection over the country, soon becoming relatively ineffective and largely unused. Figure 4, in turn, illustrates the model-generated epidemic curves of the pandemic by region (in panel (a)) as well as the model-generated cumulative curve of the epidemic as it would have occurred *without* any available vaccines (see panel (b)).

The vaccine doses made available by the CDC certainly reduced the total number of infections in the United States (see Table 2). The vaccine program was particularly effective in New England, with a 36% improvement in the total number of infections, and was quite ineffective in the Southwest, with only a 9% improvement. In Figure 5, we take a closer look at the epidemic curves and allocation strategies if we use four of the heuristics

Figure 4. Model-Generated Epidemic Curves by Region (a) and Model-Generated Cumulative Epidemic Curves With and Without Vaccine (b)



Note. Panel (b) shows the U.S. combined epidemic curve and the model-generated curve of what would have occurred if no vaccine doses were administered.

Table 2. Model-Generated Total Infection Numbers by Region With and Without Administered Vaccines

Region	Total infections		Reduction of infections (%)
	Vaccines as administered	No vaccines	
1. New England	4,464,000	6,988,000	36.12
2. New York Area	5,454,000	6,821,000	20.04
3. Mid-Atlantic	10,311,000	11,627,000	11.32
4. Southeast	13,414,000	15,218,000	11.85
5. Great Lakes	11,116,000	14,464,000	23.15
6. Southwest	8,577,000	9,488,000	9.60
7. The Plains	3,659,000	4,238,000	13.66
8. Rocky Mountains	2,414,000	2,802,000	13.85
9. West	7,002,000	9,866,000	29.03
10. Pacific Northwest	2,385,000	2,878,000	17.13
Total	68,795,000	84,391,000	18.48

discussed above: the pro rata, prepeak, critical period, and telescope-to-the-future switching heuristics. Table 3 summarizes the total number of infections under each heuristic.

The three heuristics suggested in the previous section performed better, on average, than the currently used pro rata heuristic in mitigating the total effects of the outbreak. The critical period heuristic is the least effective of the three because this is an all-or-nothing heuristic, and it depends very strongly on the correctness of the approximation algorithm that determines whether a region has entered its critical period. With the limited information available during the first few decision points, this heuristic often made entirely incorrect decisions. The best-performing heuristic was the telescope-to-the-future switching algorithm, preventing roughly 20 million infections compared with the pro rata heuristic’s 15.5 million. This is a 31% improvement in averted infections over the currently used allocation algorithm. We take a closer look at the two heuristics in Figure 6.

The prepeak heuristic, in which we gave preference to those states that were deemed to have not yet reached the peak of the epidemic, also performed well, with a 19% improvement over the pro rata heuristic. However, this all-or-nothing approach is difficult to justify and implement with uncertain data in an actual real-life scenario,

Figure 5. Epidemic Curves by Region (Top) and Decisions Made (Bottom) by Using the Pro Rata, Prepeak, Critical Period, and Telescope-to-the-Future Switching Heuristics

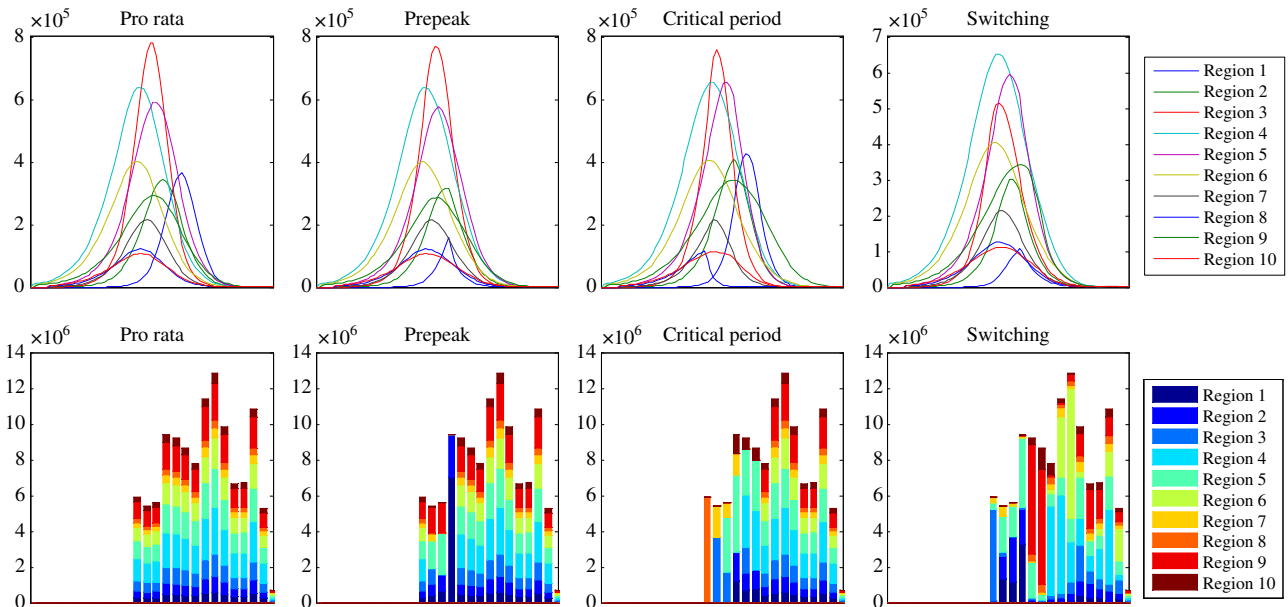


Table 3. Cumulative Number of Infections Using the Parameters Fit from the 2009–2010 H1N1 Epidemic Data

Strategy	Total infections	Difference in infections averted (%)
No vaccine	84,391,000	—
Pro rata	68,795,000	—
Prepeak	65,798,000	19.22
Critical period	67,536,000	8.07
Switching algorithm	64,012,000	30.67

especially with the potential pitfalls evident with the critical period heuristic. The telescope-to-the-future switching algorithm allows for a more robust distribution of vaccines while maintaining high effectiveness.

Improving Results

Earlier Vaccine Shipment. We must keep in mind that the numbers in Table 3 are a result of the modeling based on the 2009–2010 H1N1 scenario, in which the vaccines arrived very late and most of the states received their first vaccines after the peak of the outbreak had already occurred. It is possible that when the next deadly outbreak occurs, it will originate overseas, in Asia or even Africa, and the CDC will have more time to develop and distribute influenza vaccine doses (Diamond and Wolfe 2008). A switch from the egg-based vaccine technology currently employed to the animal-cell research used in some European countries would ensure an even shorter manufacturing time to create and ship vaccines. Such new technology could produce vaccines about 8 to 10 weeks faster than the egg-based methods (Maugh 2011, Rappuoli 2006). To examine the relative effect of timing on vaccine effectiveness, we consider what would happen if the vaccines became available two or even four weeks earlier than was the case in 2009 (see Tables 4 and 5).

In contrast to the 2009 historical experience, the number of people infected using the pro rata heuristic dropped 11 million with a two-week headway and 25 million with a four-week lead. In the case of the telescope-to-the-future switching heuristic, the additional number of infections averted—assuming that this heuristic had been used before in 2009—was also 11 and 25 million, respectively. The extra time significantly increases the effectiveness of available vaccines. And with earlier administration, the difference between pro rata and switching heuristics becomes less pronounced as timing becomes less critical. To generalize this notion, we do a similar analysis with vaccines becoming available four and eight weeks earlier, with the graphical results in Figure 7. Figures 8(a)–8(c) show the general trend of vaccine effectiveness at different shipping times.

As we see from Figure 8(b), the absolute difference between the number of averted infections with the switching algorithm and the pro rata heuristic is not monotonic. The relative difference given in percentages in Figure 8(c) is highest in the most realistic situation—when vaccines arrive relatively late, as they had in 2009. It is at that point that exploiting the real-time, dynamically evolving differences between regions becomes most important and effective, providing as much as a 31% improvement in lowering the toll of an outbreak.

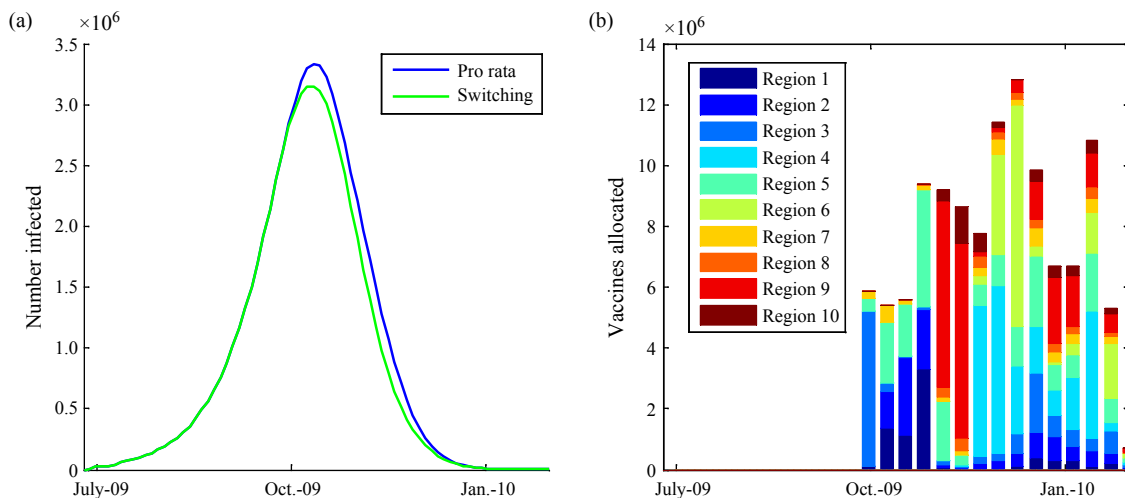
Figure 6. Cumulative Epidemic Curve for the Pro Rata and Switching Heuristic (a) and the Decisions Made Through the Telescope-to-the-Future Switching Heuristic (b)

Table 4. Cumulative Number of Infections if Vaccines Were Shipped Two Weeks Earlier Than They Were in 2009

Strategy (2 weeks earlier)	Total infections	Contrast to pro rata, difference in infections averted (%)
No vaccine	84,391,000	—
Pro rata	57,349,000	—
Prepeak	54,817,000	9.34
Critical period	55,748,000	5.9
Switching algorithm	52,628,000	17.44

Table 5. Cumulative Number of Infections if Vaccines Were Shipped Four Weeks Earlier Than They Were in 2009

Strategy (4 weeks earlier)	Total infections	Contrast to pro rata, difference in infections averted (%)
No vaccine	84,391,000	—
Pro rata	43,379,000	—
Prepeak	41,489,000	4.61
Critical period	42,560,000	2
Switching algorithm	39,776,000	8.79

Finally, we would like to discuss a point that is unlikely to occur in a pandemic scenario but may come up during a seasonal influenza outbreak, for which authorities usually have more time to prepare. When vaccines arrive particularly early—that is, when the vaccines take on a more preventative role—we might have to make a decision before having reasonable information about the epidemic curve in some regions. In particular, in a hypothetical case in which vaccines arrived 10 weeks earlier than they actually had, none of the regions would have experienced significant epidemiological activity. When faced with this situation, the switching algorithm tends to allocate very little vaccine to a region without pronounced influenza activity. When estimating parameters for a late-start region with a minimal number of infections, we estimated an approximate value for R_0 from

Figure 7. Cumulative Epidemic Curves of the Pro Rata and Switching Heuristics with Vaccines Arriving Zero, Four, and Eight Weeks Earlier Than They Did in 2009


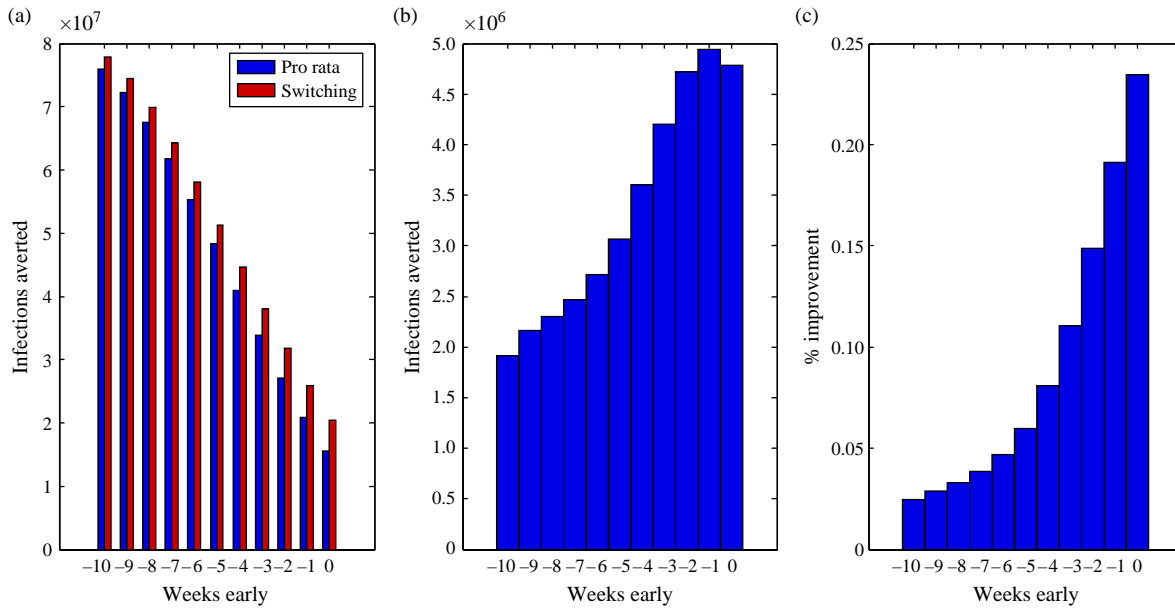


Figure 8. Total Numbers of Infections Averted as a Function of How Early the Vaccines Were Shipped (a), Difference in the Total Number of Infections Averted by the Pro Rata and Switching Heuristics (b), and Percent Difference in the Total Number of Infections Averted by These Two Heuristics (c)



other regions and assigned it to the late-start region with the assumption that the first infection will occur two weeks after the current decision time. The resulting allocation favored the regions where the outbreak was more prevalent. Of course, in the limit of months of time available before the flu wave hits, virtually any reasonable heuristic would perform well, since a large proportion of the cooperating population would be immune—thanks to vaccinations—prior to the start of any flu activity.

Higher Vaccine Uptake. The amount of vaccine administered to individuals, especially toward the end of the H1N1 pandemic, was fairly low. The total percentage was never higher than 60% of available vaccine, even in the first few days of vaccine shipment. As a last example, we provide in Table 6 the simulation results of what would happen if the total percentage of vaccine used in each region were increased by 20%.

The results encourage the use of public awareness campaigns to increase the total number of vaccines accepted by the public. A 20% increase in the acceptance of vaccines can avert anywhere from 2.3 to 7 million extra infections in the United States.

Incorporating Complexity and Constraints

Equity Constraints

Whereas the telescope-to-the-future switching heuristic might be mathematically justified by a net decrease in the total number of infections, a policy-setting body such as the CDC might feel that it is inappropriate to allocate with such a tangible real-time preference of one region over another. We can easily include this consideration in the telescope-to-the-future switching heuristic by adjusting our problem into a constrained one. We impose the constraint that we cannot take away more vaccines from a region than are regulated by certain equity limitations.

Table 6. Cumulative Number of Infections if There Was a 20% Increase in Compliance

Strategy	Total infections	Difference in infections averted (%)
No vaccine	84,391,000	—
Pro rata	66,428,000	—
Prepeak	64,085,000	13.04
Critical period	65,802,000	3.48
Switching algorithm	61,793,000	25.80

Hard Constraints. One simple way of incorporating equity is by mandating that on each day, each region i must receive no less than $l_i(t)$ vaccines; $l_i(t)$ may be an absolute number or a percentage of total available vaccines on day t .

Relative Constraints. Instead of imposing hard constraints, we might simply state that one region may not have any significant preference over another. That is, we can impose a lower-bound parameter $\rho < 1$ such that for any two regions i and j , $v_i(t)/n_i \geq \rho(v_j(t)/n_j)$.

The key to understanding how to impose these constraints is to recall that states where vaccines are inherently less effective—that is, states where the epidemic is in a late stage—are also the states where the demand for vaccines is relatively low. In those states most of the vaccines allocated in 2009 went unused. The states that received their vaccines especially late, more than five weeks after the peak of infection, used only about 15% of the vaccines allotted to them. Those states, coincidentally, were the states to which the switching algorithm allocated the fewest numbers of vaccines. Constraints based on demand forecasts will be most effective in maintaining true equity while increasing the effect of limited available vaccines.

Risk Parameter

Imposing some kind of equity constraints on the vaccine allocation problem lands us somewhere on the spectrum between the unrestricted solution obtained by the unconstrained switching algorithm and the pro rata heuristic obtained by the CDC’s actions in 2009. The place on the spectrum can be determined by the policymakers as they tailor their strategy to their constituency. We propose introducing a risk parameter, σ , that determines how nonstandard (switching) or standard (pro rata) to make our allocation strategy.

Given an allocation decision to be made on day t , if $v_i(t)$ is the decision made by the pro rata heuristic in region i on day t and $w_i(t)$ is the corresponding decision made using the switching algorithm, we allocate $(1 - \sigma)v_i(t) + \sigma w_i(t)$ to region i . The corresponding allocation lies on the spectrum of standard and nonstandard distributions. See Figure 9 for the effectiveness of the vaccines as a function of σ .

Note that the number of infections occurring during an epidemic is convex in σ . That is, the benefits of a riskier switching algorithm heuristic can be reaped with a fairly low σ value. Even a little deviation from the pro rata heuristic provides success in averting infections overall.

The benefit of using a risk parameter is that it helps avoid the pitfalls of incomplete information available in real time. A bad fit to data may cause a drastic decision to send a majority of the stockpiles to one region that might in reality not be justified. A low value of σ will safeguard against such pitfalls but will likely result in a lower number of averted infections. To make σ more adaptive, we can make it dependent on our confidence in the available data. One possible way of doing this is to create a parameter $\sigma(t)$ that is dependent on the quality of the least-squares fit to our data. We executed the simulation using $\sigma(t) = K \cdot 1/\gamma(t)$, where $\gamma(t)$ is the square norm of our fit to the existing data and K is a normalization constant. Figure 10 includes the number of infections allowed using this adaptive risk parameter.

Figure 9. Total Number of Infections Using Risk Parameter σ Combined with the Pro Rata and Switching Heuristics





Figure 10. Total Number of Infections for Various Values of σ

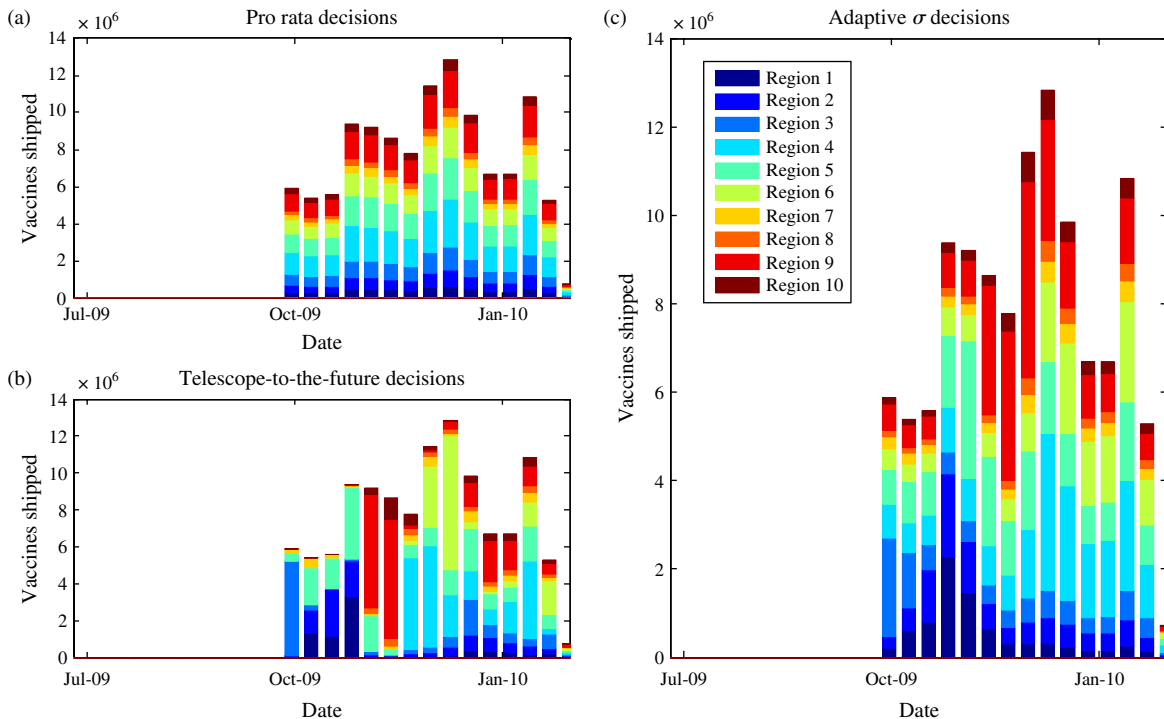


Note. The horizontal line corresponds to the threshold number of infections attained by using the adaptive risk parameter.

The adaptive risk parameter performs quite well, and the number of infections lies closer to those of the switching algorithm than the pro rata heuristic. At the same time, we can see from Figure 11 that the algorithm that uses the adaptive risk parameter allocates vaccines similarly to the pro rata heuristic at the beginning of the epidemic when data are still uncertain. This changes as time goes on and the fit to data becomes more certain.

The risk parameter can be made more complicated still by adding our forecast of vaccine uptake percentages and other uncertainty parameters; it provides a balance to prevent drastic decisions from being made. Once

Figure 11. Decisions Made by the Pro Rata (a), Telescope-to-the-Future Switching (b), and Adaptive Risk Parameters (c) Heuristics



again, we refer the reader to Teytelman (2012) for the formalized alterations to the algorithm with the addition of new constraints.

Future Work

In our development of the telescope-to-the-future algorithm, we have made an important assumption that all regions are independent of each other. That is, we assumed that the effect of vaccines allocated to one region do not affect the progression of influenza in any of the other regions under consideration. This is an oversimplification of a real-life system. Indeed, it is likely that vaccines allocated to one region will have some secondary beneficial effects on nearby regions. We propose that future research be done in representing each region as a node in a network, where edges can represent interregion proximity. Consideration of these secondary effects has the potential to increase the effectiveness of the telescope-to-the-future algorithm.

Conclusion

We have addressed the problem of allocating vaccine doses to multiple regions during a flu epidemic when vaccines are delivered late and in limited quantities over time. Our results show that changing vaccine deployments from the standard pro rata method (based on populations alone) to an adaptive strategy that takes into account the current status of the flu in each region may reduce flu infections by up to 31%. Implementation of such a strategy would require real-time use of simple and well-tested mathematical models of flu progression, fed with parameter estimates from flu incidence data sent from the field. In contrast to the pro rata policy, our suggested procedures are most effective when vaccines are delivered late. The procedures' relative advantage disappears when vaccines are available well before flu activity appears in the region.

This problem needs to be considered with the full weight of the requirements placed upon decision makers in the case of a real influenza emergency. Organizations that distribute vaccine, such as the CDC, are faced with uncertain data and epidemiological as well as political considerations that must be made every time there is vaccine to be distributed. Each decision lies on a spectrum, where less standard decisions—ones that disproportionately favor certain regions over others—result in fewer infections than do more standard decisions such as the pro rata heuristic currently used by the CDC. We provide in this work the methodology that can help gain insights into the effectiveness of vaccines at different times in the outbreak, and we evaluate how well various deployment strategies perform in reducing the total number of infections. This methodology is a tool in the hands of the decision maker, who must choose just how standard or nonstandard the allocation method should be. With that in mind, a combination of the switching heuristic, the pro rata heuristic, and perhaps other allocation schemes could provide the most epidemiologically effective and politically viable strategy for mitigating the toll of influenza.

Acknowledgments

Work on this paper was supported under a cooperative agreement with the U.S. Centers for Disease Control and Prevention (CDC) [Grant 1 PO1 TP000307-01, “LAMPS (Linking Assessment and Measurement to Performance in PHEP Systems)”] awarded to the Harvard School of Public Health Center for Public Health Preparedness (HSPHCPHP) and the Massachusetts Institute of Technology (MIT) Center for Engineering Systems Fundamentals (CESF), and by the Sloan Foundation of New York [Grant 2007-3-11, “Decision-Oriented Analysis of Pandemic Flu Preparedness and Response”]. The discussion and conclusions in this paper are those of the authors and do not necessarily represent the views of the CDC, the Sloan Foundation, the U.S. Department of Health and Human Services, Harvard, or MIT. The authors thank Stan Finkelstein for helpful comments on an earlier draft.

Appendix A. Efficient Switching Algorithm

Initialize $v_i(k) = v_i^*(k)$ for $k < t$ and $v_i(k) = v(k) \cdot (n_i / \sum_{i=1}^R n_i)$ for $k \geq t$ for all i .

Initialize a $R \times R \times (T - t + 1)$ matrix A .

Let $\delta_v = 2^{\lfloor \log(\max_{i,k} v_i(k)) \rfloor}$

$\delta_c = 2^{\lfloor \log(I(v^*)) \rfloor}$

while $\delta_c \geq 1$

$\delta_v = 2^{\lfloor \log(\max_{i,k} v_i(k)) \rfloor}$

while $\delta_v \geq \delta_c$

while $(\exists i, j, k \mid I(\vec{v}_i) + I(\vec{v}_j) > I(\vec{v}_i + \delta_v \cdot \vec{e}_k) + I(\vec{v}_j - \delta_v \cdot \vec{e}_k))$

```

for every triple  $(i, j, k)$  such that  $\vec{v}_j(k) > \delta_v$ 
  calculate  $A(i, j, k) = I(\vec{v}_i + \delta \cdot \vec{e}_k) + I(\vec{v}_j - \delta \cdot \vec{e}_k)$ 
  let  $(i^*, j^*, k^*) = \arg \min(A)$ 
  if  $A(i^*, j^*, k^*) > \delta_c$ 
     $\vec{v}_{i^*} = \vec{v}_{i^*} + \delta_v \cdot \vec{e}_{k^*}$ 
     $\vec{v}_{j^*} = \vec{v}_{j^*} - \delta_v \cdot \vec{e}_{k^*}$ 
  else
    if  $A(i^*, j^*, k^*) > 0$ 
       $\vec{v}_{i^*} = \vec{v}_{i^*} + \delta_v \cdot \vec{e}_{k^*}$ 
       $\vec{v}_{j^*} = \vec{v}_{j^*} - \delta_v \cdot \vec{e}_{k^*}$ 
    end
     $\delta_v = \delta_v/2$ 
  end
end
end
end
end
 $\delta_c = \delta_c/2$ 
end

```

This algorithm is a modified version of the gradient descent algorithm resulting from the inherently discrete nature of the problem. We start with moving large numbers of vaccines but only when the cost improvement is significantly large. If moving large numbers of vaccines is not possible, we lower the cost and vaccine thresholds until no vaccines can be moved to any improvement at all.

Appendix B. Stochastic Spread Model

We used the deterministic heterogeneous community influenza spread model (Teytelman and Larson 2012) as a basis for the stochastic model. The model had two levels of heterogeneity, λ and s . At the onset of the epidemic, λ was distributed as a Poisson distribution with mean μ over the entire population, μ was found by using a fit to the data from the 2009–2010 H1N1 epidemic, and s corresponded to an individual's susceptibility to infection given contact with an infectious person, set to have two values: one for a nonvaccinated person and one for a person that has been vaccinated against the current influenza strain. We modify the deterministic model to incorporate the stochastic nature of a real outbreak during our Monte Carlo simulations.

Recall from the deterministic heterogeneous community influenza spread model that for each day i we calculate the following quantities:

$$\begin{aligned}
 N_{i+1}^I &= N_i^S \int_0^1 \int_0^\infty p(I | \lambda, s) \cdot f_i^S(\lambda, s) d\lambda ds, \\
 N_{i+1}^S &= N_i^S - N_{i+1}^I, \\
 f_{i+1}^I(\lambda, s) &= \frac{p_i(I | \lambda, s) \cdot f_i^S(\lambda, s)}{\int_0^1 \int_0^\infty p(I | \lambda, q, s) \cdot f_i^S(\lambda, q, r, s) d\lambda ds}, \\
 f_{i+1}^S(\lambda, s) &= \frac{(1 - p_i(I | \lambda, s)) \cdot f_i^S(\lambda, s)}{\int_0^1 \int_0^\infty (1 - p(I | \lambda, s)) \cdot f_i^S(\lambda, s) d\lambda ds}.
 \end{aligned}$$

When calculating the same quantities for a stochastic model, we follow similar steps. The probability that a given interaction is with an infectious person remains as before:

$$\beta_i = \frac{N_i^I E_i^I(\lambda)}{NE(\lambda)}.$$

Similarly, the probability of becoming infected given an infectivity level λ ,

$$p_i(I | \lambda, s) = 1 - e^{-\lambda s \beta_i},$$

remains the same.

Thus, the number of infections of people of infectivity level λ and susceptibility level s is a random variable picked from a binomial distribution where n is the number of susceptible individuals with parameter λ and $p = p_i(I | \lambda, s) = 1 - e^{-\lambda s \beta_i}$.

After selecting a value from this binomial distribution for each pair (λ, s) , we proceed as before to calculate $N_{i+1}^I \cdot N_{i+1}^S$, $f_{i+1}^I(\lambda, s)$, and $f_{i+1}^S(\lambda, s)$.

For $|(1/\sqrt{n})(\sqrt{(1-p)/p} - \sqrt{p/(1-p)})| < 0.3$ (Box et al. 1978), the binomial random variables were approximated with the appropriate normal distribution.

Appendix C. Parameter Values

We used the 10 regions specified by the CDC and fit our parameters to the data available from 2009 to 2010. Table C.1 contains the calculated percentage of the population covered by administered vaccine by region, and Table C.2 replicates the official vaccine numbers administered and shipped to the 10 U.S. regions. In Tables C.3 and C.4, we provide fitted parameters used in our model, and the resulting estimated epidemic curves in each region are provided in Figure C.1.

Table C.1. Vaccine Acceptance Rate in U.S. Regions

Region	Population	% covered by administered vaccine
1	14,429,720	32.6
2	19,949,192	20.7
3	28,891,734	23
4	60,580,377	20.1
5	51,766,882	24.2
6	37,860,549	17.8
7	13,610,802	23.5
8	10,787,806	25.4
9	47,495,705	22.1
10	12,734,126	25

Table C.2. Vaccine Schedule (Centers for Disease Control and Prevention 2010a)

Date	Number allocated	Numbered ordered	Number shipped
01/29/10	147,323,810	119,086,300	118,922,220
01/28/10	147,301,010	118,989,300	118,899,820
01/27/10	147,089,010	118,920,900	118,841,120
01/26/10	146,940,510	118,840,000	118,606,720
01/25/10	146,267,090	118,699,900	118,279,120
01/22/10	145,132,850	118,630,900	116,226,920
01/21/10	144,741,750	118,072,700	116,206,320
01/20/10	144,092,400	117,711,000	115,822,020
01/19/10	141,602,700	115,853,000	115,630,620
01/15/10	139,334,000	115,591,600	115,487,820
01/14/10	138,934,400	115,246,700	114,850,120
01/13/10	138,402,600	114,972,000	114,504,220
01/12/10	138,369,200	114,517,000	113,293,220
01/11/10	137,697,600	113,270,700	112,538,020
01/08/10	137,359,300	112,599,500	110,202,920
01/07/10	135,981,600	112,013,100	110,202,520
01/06/10	131,661,400	110,400,200	109,016,520
01/05/10	130,386,600	109,506,200	107,351,020
01/04/10	130,386,600	108,474,200	105,266,520
12/31/09	118,241,100	106,247,500	99,366,920
12/30/09	118,241,100	104,249,900	99,362,620
12/29/09	116,266,700	101,697,200	98,764,920
12/28/09	111,947,700	100,931,500	97,281,820
12/24/09	110,477,700	98,681,000	92,659,820
12/23/09	107,786,500	97,584,000	92,659,820
12/22/09	107,786,500	95,514,300	91,863,720
12/21/09	107,156,700	93,940,500	90,384,920
12/18/09	100,082,700	91,630,300	85,988,420
12/17/09	99,439,500	89,187,200	85,980,220
12/16/09	96,385,200	86,710,200	84,493,120
12/15/09	94,628,100	85,298,200	82,617,120
12/14/09	92,883,100	83,656,200	76,355,920


Table C.3. Fitted Parameters

Region	μ	R_0	Epidemic start time	Initial number infected
1	13.32	1.43	Aug 13, 2009	508
2	11.66	1.26	July 21, 2009	905
3	12.29	1.32	July 28, 2009	2,921
4	10.81	1.18	June 30, 2009	10,000
5	11.06	1.21	July 11, 2009	4,607
6	10.80	1.18	July 2, 2009	8,940
7	11.35	1.24	July 14, 2009	1,710
8	10.88	1.18	July 11, 2009	3,386
9	10.44	1.14	June 30, 2009	6,128
10	10.58	1.16	July 11, 2009	5,656

Table C.4. Other Fixed Global Parameters

Parameter	Value
$s_{\text{vaccinated}}$ (susceptibility)	0.02
$s_{\text{unvaccinated}}$ (susceptibility)	0.1
Duration of generation	2.3 calendar days
Number of iterations for each heuristic	500

Figure C.1. Approximated Epidemic Curves by Region



References

Box GEP, Hunter JS, Hunter WG (1978) *Statistics for Experimenters: Design, Innovation, and Discovery*, 2nd ed. (John Wiley & Sons, New York).

Cameiro HA, Mylonakis E (2009) Google trends: A Web-based tool for real-time surveillance of disease outbreaks. *Clinical Infect. Dis.* 49(10):1557–1564.

Centers for Disease Control and Prevention (2009) Media briefing: Update on 2009 H1N1 flu. Press release (September 3), CDC, Atlanta. <http://www.cdc.gov/media/transcripts/2009/t090903.htm>.

Centers for Disease Control and Prevention (2010a) 2009 H1N1 vaccine doses allocated, ordered and shipped by project area. Accessed January 20, 2011, <http://www.cdc.gov/h1n1flu/vaccination/vaccinesupply.htm>.

Centers for Disease Control and Prevention (2010b) Interim results: Influenza A (H1N1) 2009 monovalent and seasonal influenza vaccination coverage among health-care personnel—United States, August 2009–January 2010. *Morbidity Mortality Weekly Rep.* 59(12):357–362.

- Centers for Disease Control and Prevention (2010c) The 2009 H1N1 pandemic: Summary highlights, April 2009–April 2010. Accessed January 12, 2011, <http://www.cdc.gov/h1n1flu/cdcreponse.htm>.
- Centers for Disease Control and Prevention (2012) Selecting the viruses in the seasonal influenza (flu) vaccine. Accessed August 1, 2012, <http://www.cdc.gov/flu/professionals/vaccination/virusqa.htm>.
- Chowell G, Viboud C, Wang X, Bertozzi S, Miller M (2009) Adaptive vaccination strategies to mitigate pandemic influenza: Mexico as a case study. *PLoS One* 4(12):e8164.
- Diamond J, Wolfe N (2008) Where will the next pandemic emerge? *Discover* (November) <http://discovermagazine.com/2008/nov/27-where-will-the-next-pandemic-emerge>.
- Finkelstein SN, Hedberg KJ, Hopkins JA, Hashmi S, Larson RC (2011) Vaccine availability in the United States during the 2009 H1N1 outbreak. *Amer. J. Disaster Med.* 6(1):23–30.
- Fraser C, Donnelly CA, Cauchemez S, Hanage WP, Van Kerkhove MD, Hollingsworth TD, Griffin J, et al. (2009) Pandemic potential of a strain of influenza A (H1N1): Early findings. *Science* 324(5934):1557–1561.
- Hardelid P, Fleming DM, McMenamin J, Andrews N, Robertson C, Sebastian Pillai P, Ellis J et al. (2011) Effectiveness of pandemic and seasonal influenza vaccine in preventing pandemic influenza A(H1N1)2009 infection in England and Scotland 2009–2010. *Eurosurveillance* 16(2):Article 3.
- Harder KM, Andersen PH, Baehr I, Nielsen LP, Ethelberg S, Glismann S, Mølbak K (2011) Electronic real-time surveillance for influenza-like illness: Experience from the 2009 influenza A(H1N1) pandemic in Denmark. *Eurosurveillance* 16(3):Article 1.
- Harris KM, Maurer J, Kellerman AL (2010) Influenza vaccine—Safe, effective and mistrusted. *New Engl. J. Med.* 363(23):2183–2185.
- Hill AN, Longini IM (2003) The critical vaccination fraction for heterogeneous epidemic models. *Math. Biosciences* 181(1):85–106.
- Hopkins J (2011) H1N1 after-action reports: Lessons on vaccine distribution. ESD working paper, Massachusetts Institute of Technology, Cambridge. <http://esd.mit.edu/WPS/2011/esd-wp-2011-06.pdf>.
- Larson RC, Teytelman A (2012) Modeling the effects of H1N1 influenza vaccine distribution in the United States. *Value Health* 15(1):158–166.
- Longini IM Jr, Halloran ME (2005) Strategy for distribution of influenza vaccine to high-risk groups and children. *Amer. J. Epidemiol.* 161(4):303–306.
- Longini IM Jr, Nizam A, Xu S, Ungchusak K, Hanshaoworakul W, Cummings DA, Halloran ME (2005) Containing pandemic influenza at the source. *Science* 309(5737):1083–1087.
- Maugh TH II (2011) Cell-culture influenza vaccine proves effective, could speed production. *Los Angeles Times* (February 16) <http://articles.latimes.com/2011/feb/16/news/la-heb-influenza-vaccine-02162011>.
- Nigmatulina KR, Larson RC (2009) Living with influenza: Impacts of government imposed and voluntarily selected interventions. *Eur. J. Oper. Res.* 195(2):613–627.
- Patel R, Longini IM Jr, Halloran ME (2005) Finding optimal vaccination strategies for pandemic influenza using genetic algorithms. *J. Theoret. Biol.* 234(2):201–212.
- Rappuoli R (2006) Cell-culture-based vaccine production: Technological options. *Bridge* 36(3):25–30.
- Teytelman A (2012) Modeling reduction of pandemic using influenza spread pharmaceutical and nonpharmaceutical interventions in a heterogeneous population. Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge.
- Teytelman A, Larson RC (2012) Modeling influenza progression within a continuous-attribute heterogeneous population. *Eur. J. Oper. Res.* 220(1):238–250.