European Journal of Operational Research 233 (2014) 208–219

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Optimal two-phase vaccine allocation to geographically different regions under uncertainty

Hamed Yarmand [a,*], Julie S. Ivy [a], Brian Denton [c], Alun L. Lloyd [b]

[a] Department of Industrial and Systems Engineering, North Carolina State University, Raleigh, NC, USA
[b] Department of Mathematics, North Carolina State University, Raleigh, NC, USA
[c] Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI, USA

ARTICLE INFO

Article history:
Received 27 February 2013
Accepted 19 August 2013
Available online 31 August 2013

Keywords:
OR in health services
Epidemic control
Two-phase vaccine allocation
Stochastic linear program
Newsvendor model
Value of stochastic solution

ABSTRACT

In this article, we consider a decision process in which vaccination is performed in two phases to contain the outbreak of an infectious disease in a set of geographic regions. In the first phase, a limited number of vaccine doses are allocated to each region; in the second phase, additional doses may be allocated to regions in which the epidemic has not been contained. We develop a simulation model to capture the epidemic dynamics in each region for different vaccination levels. We formulate the vaccine allocation problem as a two-stage stochastic linear program (2-SLP) and use the special problem structure to reduce it to a linear program with a similar size to that of the first stage problem. We also present a Newsvendor model formulation of the problem which provides a closed form solution for the optimal allocation. We construct test cases motivated by vaccine planning for seasonal influenza in the state of North Carolina. Using the 2-SLP formulation, we estimate the value of the stochastic solution and the expected value of perfect information. We also propose and test an easy to implement heuristic for vaccine allocation. We show that our proposed two-phase vaccination policy potentially results in a lower attack rate and a considerable saving in vaccine production and administration cost.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

In case of an outbreak of an infectious disease, such as influenza, one of the most effective interventions is vaccinating the susceptible population. Vaccination is known to be more effective at the beginning of the epidemic (Khazeni, Hutton, Garber, Hupert, & Owens, 2009; Yarmand, 2010; Yarmand & Ivy, 2013; Yarmand & Ivy, in press; Yarmand, Ivy, & Roberts, in press). However, it is not always possible to vaccinate a large population in a short time due to insufficient available vaccine doses or limited capacity. For instance, during the 2010–2011 flu season in the United States influenza vaccine doses were administered continuously from August 2010 through May 2011 (Centers for Disease Control and Prevention (CDC), 2011a). The national vaccination coverage was estimated at 30% by the end of October and 43% by the end of May (the situation was similar for flu season 2009–2010) (CDC, 2011a). These observations suggest a new vaccination policy: to have two (or more) phases of vaccination in geographically different regions when the first phase occurs at the beginning of the epidemic (to have the maximum effect) and the vaccination level in the second phase in each region depends on the outcome of vaccination in the first phase in that region. Under this policy, which we refer to as the "two-phase vaccination policy", one critical question is "how many vaccine doses should be allocated to each region in each phase?".

The main advantage of the two-phase vaccination policy is to allow evaluation of the vaccination outcome after the first phase. Therefore phase two might not be necessary in some regions resulting in a potential decrease in vaccine production as well as possibility of redistribution of vaccine doses among regions which need the second phase. Note that this two-phase model is a special case of a more general multi-phase model. However having two phases is a reasonable assumption since it is unlikely such allocation decisions would be made on a frequent basis. Furthermore, evaluating the two-phase model provides a conservative estimate of the benefits of a more general multi-phase model.

In the United States, vaccine stockpiles in each state are centralized and controlled by the state health department. Under a two-phase vaccination policy the state health department would need to decide how many vaccine doses to allocate to each region (e.g., each county) in each phase, and hence how many vaccine doses to order for production. The allocation in the first phase is challenging because of uncertainty about the epidemic dynamics. The epidemic may progress, or die out in some regions with limited vaccination, or even without any vaccination. Due to uncertainty, we formulate the problem of allocating vaccine doses to different

* Corresponding author. Address: Francis H. Burr Proton Therapy Center, Department of Radiation Oncology, MGH, Harvard Medical School, 30 Fruit St., Boston, MA 02114, USA. Tel.: +1 617 724 3665; fax: +1 919 515 5281.
E-mail addresses: hyarman@ncsu.edu (H. Yarmand), jsivy@ncsu.edu (J.S. Ivy), btdenton@umich.edu (B. Denton), alun_lloyd@ncsu.edu (A.L. Lloyd).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.08.027

regions in different phases of vaccination as a stochastic resource allocation problem.

In this article, we consider seasonal influenza in geographically different regions (e.g., counties within a state). We present numerical results based on the 100 counties of the state of North Carolina. We examine the two-phase vaccination policy for a range of choices of model input parameters. We assume that health care officials have decided to have vaccination in two distinct time periods referred to as Phase-I and Phase-II. Based on the probability that the epidemic is contained in Phase-I in each region (estimated from the disease spread model discussed in Section 3.1), we formulate a two-stage stochastic linear program (2-SLP). The Phase-I vaccine allocation is determined in the first stage while the second stage determines Phase-II vaccine allocation according to the realizations of epidemic containment in different regions in Phase-I, hence providing flexibility in the Phase-II vaccine allocation. We use the special problem structure to reduce the developed 2-SLP to a linear program (LP) with a similar size to that of the first stage problem and therefore find the optimal allocation very efficiently.

We use our model for three purposes: (i) Estimating the saving in vaccine doses and the associated cost as a result of implementing the two-phase vaccination policy compared to the current continuous vaccination policy, (ii) estimating the optimal number of vaccine doses that should be ordered for production prior to the flu season, and (iii) estimating the associated value of the stochastic solution (*VSS*) and the expected value of perfect information (*EVPI*). The *VSS* provides an estimate of the usefulness of considering the stochastic nature of the problem while the *EVPI* provides an estimate of the value of epidemic detection and forecast systems. We also propose a heuristic for vaccine allocation based on our numerical results.

In addition to the 2-SLP formulation, we will also present a Newsvendor formulation of this problem. The Newsvendor formulation allows the identification of a closed form solution for the optimal allocation. It also provides insight into the structure of the optimal allocation that is not provided by the 2-SLP model; however, as we show this comes at the expense of requiring the complete distribution of demand, as opposed to just the mean demand for the 2-SLP mode (after it is reduced to an LP).

The remainder of this article is organized as follows. In Section 2 we present a review of relevant research. In Section 3 we present the problem definition and assumptions and present the 2-SLP and Newsvendor model formulations along with our solution methodology. In Section 4 we present and analyze the numerical results. Finally in Section 5 we conclude this article with summarizing the main findings and suggesting some directions for future research.

2. Literature review

Several aspects of flu vaccination have been analyzed by researchers including vaccine strain selection (e.g., Cho (2010), Kornish and Keeney (2008), and Wu, Wein, and Perelson (2005)), vaccine supply chain (e.g., Chick, Mamani, and Simchi-Levi (2008) and Deo and Corbett (2008)), vaccine market and the associated economic impacts (e.g., Brito, Sheshinski, and Intriligator (1991), Geoffard and Philipson (1996), and Philipson (2000, chap. 33)), and vaccination decisions among individuals (e.g., Bauch and Earn (2004), Chapman and Coups (1999), Galvani, Reluga, and Chapman (2007), Larson, Olsen, Cole, and Shortell (1979), and Reluga, Bauch, and Galvani (2006)) and households (e.g., Yarmand and Ivy (in press)).

In the context of vaccine supply chain management, the Newsvendor model has sometimes been used. For example, Chick et al. (2008) use a Newsvendor model for analyzing the influenza vaccine supply chain. In their Newsvendor model, the demand uncertainty is replaced by production uncertainty. Therefore they used the Newsvendor model in the production stage, and not the allocation stage.

In the context of epidemic control, a number of stochastic resource allocation models have been developed to determine optimal vaccination strategies. Becker and Starczak (1997) study vaccination policies in a stochastic SIR model (a model with susceptible, infective, and recovered population; see Hethcote (2000) for details) when the population has been divided into a community of households. They derive a closed form equation for the post-vaccination reproduction number, R_* (first introduced by Ball, Mollison, and Scalia-Tomba (1997)). Also they consider the constraint $R_* \leqslant 1$ and formulate and numerically solve an LP to find the minimum vaccination coverage under this constraint which ensures that the epidemic will die out.

Tanner, Sattenspiel, and Ntaimo (2008) present a stochastic programming framework for finding the optimal vaccination policy for controlling infectious disease epidemics under parameter uncertainty. They initially present a model to find the vaccination policy with the minimum cost under a chance constraint. The chance constraint requires $\Pr(R_* \leqslant 1) \geqslant \alpha$, where α is a predetermined parameter defined by the decision maker. They also present the problem formulation for two additional cases: (a) finding the optimal vaccination policy when vaccine supply is limited and (b) a cost-benefit scenario. They extend the LP formulation of Becker and Starczak (1997) to this stochastic programming framework.

Drawing upon the earlier work of Ball et al. (1997), Ball and Lyne (2002) consider the case of an all-or-nothing vaccine (i.e., a person is either completely immune following vaccination or the vaccine is completely ineffective). They show that if the sequence $\{n\mu_n\}$ is convex, where μ_n is the mean size of a local outbreak within a household of size n, then the optimal solution to the LP problem formulation of Becker and Starczak (1997) can be characterized explicitly.

Finally, Hill and Longini (2003) use a general framework that could apply to several epidemic situations including incorporation of latent periods (resulting in the SEIR model which also includes the exposed population; see Hethcote (2000) for details), diseases with permanent immunity (SIR model), or no immunity (SIS model in which infectives become susceptible again after recovery) with and without vital dynamics. They develop a method to derive optimal vaccination strategies for populations divided into heterogeneous subgroups.

We note that all of these models have considered only one phase of vaccination as opposed to two phases in our model. Therefore they do not allow evaluation of the vaccination outcome in the middle of the epidemic. Only a few of these models have considered vaccination cost in their analysis while our objective is to minimize the expected number of administered vaccine doses, and hence the expected cost of vaccination. In addition, most of these models have considered uncertainty in the disease parameters, while we have considered uncertainty in the vaccination outcome regarding the epidemic control which accounts for a different aspect of the stochastic nature of the problem in the context of resource allocation decisions. Finally, we will present a Newsvendor formulation of the vaccine allocation problem which, to the best of our knowledge, is the first application of the Newsvendor model for vaccine allocation.

3. Problem definition and assumptions

We assume that there are N regions to be vaccinated in an attempt to contain an epidemic (seasonal influenza in our study). Let $\mathcal{N} = \{1, 2, \ldots, N\}$ denote the set of indices of different regions

and P_i denote the population of region $i \in \mathcal{N}$. We say that the epidemic is contained in region $i \in \mathcal{N}$ if the attack rate, defined as the percentage of population who become infective during the flu season, in region i does not exceed a predetermined threshold denoted by *ART* (attack rate threshold). The attack rate for seasonal influenza (and therefore *ART*) ranges from 5% to 15% (Homeland Security, 2011; World Health Organization, 2009). In our model we assume vaccination occurs in two phases: Phase-I occurs before the flu season and Phase-II occurs in the middle of the flu season.

We suppose that a total of V_1 and V_2 doses of vaccine will be available in Phase-I and Phase-II, respectively. For seasonal influenza, vaccine production begins at least six months before the epidemic (Gerdil, 2003), and therefore both V_1 and V_2 are known at the time the epidemic begins. Since we will compare our proposed two-phase vaccination policy to the current continuous vaccination policy, we estimate V_1 and V_2 from the current vaccine distribution data. We let V_1 be the total vaccine doses distributed before the flu season begins. Let α denote the (average) vaccination percentage coverage in previous flu seasons in the considered regions (in the rest of this article "coverage" is always measured in percentages). We let $V_2 = \alpha \sum_{i=1}^{N} P_i - V_1$, which represents the maximum demand for vaccine doses in all regions in Phase-II.

We denote the number of vaccine doses allocated to region $i \in \mathcal{N}$ in Phase-I by x_i and in Phase-II by y_i as the first and second stage decision variables, respectively. In our model we assume x_i and y_i, $i \in \mathcal{N}$, are nonnegative continuous variables due to their large value.

Our model captures an important tradeoff between vaccination phases for diseases such as the seasonal influenza. In Phase-I, which occurs before the beginning of the flu season, available vaccine doses may not be sufficient for a large portion of the population, hence the possible need for Phase-II vaccination. Also by accumulating the administered vaccine doses before the flu season begins into Phase-I, the chance of containing the epidemic increases and also the attack rate decreases due to early vaccination. Furthermore, the epidemic might be contained as the result of Phase-I vaccination in some regions. Therefore no further vaccination is necessary in Phase-II in those regions, resulting in saving vaccine doses for regions which need Phase-II vaccination (or even for regions other than the considered regions) and, on average, a decrease in the total required vaccine doses, and hence, vaccine production level and the associated cost.

For a given value of *ART*, we denote by $F_i(v_i)$ the probability of epidemic containment in region $i \in \mathcal{N}$ with Phase-I vaccination coverage v_i (i.e., with $100v_i\%$ of the population vaccinated in Phase-I). The function $F_i(\cdot)$, $i \in \mathcal{N}$, represents the cumulative distribution function (CDF) of the vaccination coverage required in Phase-I to contain the epidemic in region i. We denote by ξ_i the associated binary random variable whose value is 0 if the epidemic is contained in region $i \in \mathcal{N}$ as a result of Phase-I vaccination (with probability $F_i(v_i)$) and 1 if the epidemic is not contained in region i (with probability $1 - F_i(v_i)$). Random variables ξ_i, $i \in \mathcal{N}$, are independent from each other because we consider importation of infectives (i.e., those who are capable of transmitting the disease to susceptible individuals) from other regions in the disease spread model for estimating $F_i(\cdot)$, $i \in \mathcal{N}$ (see Section 3.1). In other words, the impact that different regions have on each other due to travel of infectives has been accounted for in the process of estimating $F_i(\cdot)$, $i \in \mathcal{N}$. It should be noted that $F_i(\cdot)$, $i \in \mathcal{N}$, are estimated for planning purposes before the flu season begins. Therefore it is not possible to use observation of the vaccination outcome in one region to have a better estimate of epidemic containment probability in the neighboring regions. As a result, we have to use the parameters estimated from the data from previous flu seasons to simulate the disease spread in different regions. However, we can approximate the impact of different regions on each other by accounting for the travel of infectives to capture the dependence between regions.

As discussed previously, seasonal influenza vaccination begins before the beginning of the flu season to contain the epidemic. Accordingly, it is reasonable to assign a minimum number of vaccine doses to each region in Phase-I. Therefore, we assume health care officials assign a pre-determined minimum coverage for each region. We denote the minimum number of vaccine doses allocated to region $i \in \mathcal{N}$ in Phase-I by n_i. We assume that $\sum_{i=1}^{N} n_i \leqslant V_1$, otherwise it is impossible to allocate at least n_i vaccine doses to all regions. Further, we assume that n_i, $i \in \mathcal{N}$, are proportional to the population of different regions. This assumption is consistent with the issue of equity and also the demand for vaccination. Therefore, the minimum vaccination coverage (in percentage) in all regions will be the same. Let v_0 denote the minimum vaccination coverage in Phase-I in all regions (we will vary v_0 to find its optimal value, v_0^*). Therefore, we have $n_i = v_0 P_i$, $i \in \mathcal{N}$. As a result, the probability of epidemic containment for a given *ART* in region $i \in \mathcal{N}$ will be at least $F_i(v_0)$. We assume that the epidemic will be contained in region $i \in \mathcal{N}$ in Phase-I with probability $F_i(v_0)$ even if more than n_i vaccine doses are allocated to region i in Phase-I. This assumption results in linearity and tractability of the model. As we will see in Section 4, at the optimal coverage v_0^* the probability of epidemic containment is indeed $F_i(v_0^*)$, $i \in \mathcal{N}$.

If the epidemic is not contained in Phase-I in region $i \in \mathcal{N}$, then Phase-II vaccination will be implemented and further vaccine doses will be allocated to region i until a total number of at least m_i vaccine doses are allocated to region i in Phase-I and Phase-II (note that this implies $n_i \leqslant m_i$, $i \in \mathcal{N}$). We let $m_i = \alpha P_i$, $i \in \mathcal{N}$, which represents the maximum demand for vaccine doses in region i.

We assume that administering one dose of vaccine in region $i \in \mathcal{N}$ incurs a cost equal to c_i in Phase-I and d_i in Phase-II. For c_i, $i \in \mathcal{N}$, we consider the per-dose vaccine cost for CDC (CDC, 2012a) with an average of \$10 per dose. If for some region i' we have $c_{i'} \geqslant d_{i'}$, it is optimal to allocate only the minimum vaccine doses, $n_{i'}$, to region i' in Phase-I. Therefore in this case the optimal allocation is trivial. Accordingly, we are more interested in finding the optimal allocation of vaccine doses among regions i for which $c_i \leqslant d_i$. This is reasonable for seasonal diseases, such as seasonal influenza, because providing additional vaccine doses in Phase-II may be more costly than Phase-I for three reasons. First, there are several overhead costs in Phase-II such as shipment, storage, and reestablishment of mass vaccination clinics. Second, the additional vaccine doses in Phase-II might be provided by redistribution of vaccine doses among the regions which need Phase-II vaccination, hence an increase in the transportation cost. Third, the additional vaccine doses might be provided by outsourcing, which usually yields a higher cost. We let $d_i = (1 + r)c_i$, $i \in \mathcal{N}$, where r represents the possible percentage increase in vaccination cost from Phase-I to Phase-II. Without loss of generality, we make the following assumption by reordering different regions.

Assumption 1. Phase-II vaccination cost is nonincreasing in the regions index, i.e., $d_1 \geqslant d_2 \geqslant \cdots \geqslant d_N$.

For our numerical experiment, we have considered 100 counties of the state of North Carolina. We obtained the population data of 100 counties of North Carolina from the 2010 United States Census (U.S. Census Bureau, 2012). The smallest county is Tyrrell with population 4407 and the largest county is Mecklenburg with population 919,628. The flu season in North Carolina is October through May (North Carolina Public Health, 2012). Therefore we assume Phase-I occurs in September (before the flu season) and Phase-II occurs in early December.

As mentioned previously, distribution of vaccine doses for seasonal influenza begins in August. We let V_1 be the total vaccine doses distributed from August through October. For the flu season (2011–2012), a total of 124.9 million doses were distributed nationwide by the end of October (CDC, 2012b). Assuming vaccine doses were distributed among different states in proportion to their population, North Carolina, with a total population of 9,535,483 accounting for 3.09% of the United States population (U.S. Census Bureau, 2012), had a share of $V_1 = 3{,}857{,}486$ doses. For North Carolina $\alpha = 0.45$ (CDC, 2010, 2011c) and therefore $V_2 = 433{,}482$ doses. Table 1 summarizes the model parameters and variables.

3.1. Disease spread model

An important input of the model is the probability of epidemic containment with Phase-I vaccination, i.e., $F_i(v_0),\ i \in \mathcal{N}$, for specific values of *ART* and v_0. To estimate $F_i(v_0),\ i \in \mathcal{N}$, we developed a disease spread simulation model based on the stochastic SEIR epidemic model (see Hethcote (2000) and Li, Graef, Wand, and Karsai (1999) for a detailed discussion of this model). This model is appropriate for the spread of influenza and has been widely used in the literature (e.g., see Chowell, Nishiura, and Bettencourt (2007) and Flahault, Vergu, Coudeville, and Grais (2006)). We ran the simulation model many times (see Section 4 for details) and estimated $F_i(v_0),\ i \in \mathcal{N}$, based on the proportion of replications in which the calculated attack rate, defined as the percentage of individuals who have ever become infective during the simulation horizon, did not exceed *ART*.

Assuming a coverage of v_0 in Phase-I, simulation of the disease spread in region i commences with $v_0 P_i$ initial vaccinated individuals. We varied v_0 from 0% to 45%, which is the vaccination coverage for seasonal influenza observed in North Carolina in previous flu seasons (CDC, 2010, 2011c). We considered 60% effectiveness for seasonal influenza vaccine (CDC, 2011b) and accordingly assumed that a vaccinated individual becomes completely immune to the disease with probability 0.6 and remains susceptible to the disease with probability 0.4. Following a conservative approach, we assume no cross immunity in the population as a result of vaccination or infection in previous flu seasons—this is also consistent with CDC guidelines (CDC, 2011e). At the beginning of each run of the simulation model the number of individuals who have been effectively vaccinated is randomly generated based on the vaccine effectiveness.

The basic reproduction number, often denoted by R_0, for seasonal influenza is estimated at $R_0 = 1.3$ for the United States (Chowell, Miller, & Viboud, 2008). The exposure rate (i.e., rate at which susceptibles become exposed), also known as the horizontal incidence, is then calculated accordingly (see Hethcote (2000) and Li et al. (1999) for details). We considered two days as the average latent period (i.e., the period between exposure and becoming infective) and seven days as the average infectious period (i.e., the period between becoming infective and recovery) (CDC, 2011d). The infection rate (i.e., rate at which exposed individuals become infective) and the recovery rate (i.e., rate at which infectives recover) are then calculated accordingly.

At each step of the simulation model, the exposure, infection, and recovery rates are calculated. Then the time and the type of the next event (exposure, infection, or recovery) are determined. The simulation time is then advanced and the number of susceptible, exposed, infective, and recovered individuals are updated.

The disease spread model simulates the spread of influenza from the beginning of October until the end of the flu season. We used the weekly percentage of patient visits for influenza-like illness (ILI) in North Carolina in the 2010–2011 flu season (North Carolina Public Health, 2011) to calibrate the disease spread model and set the value of the associated model parameters. We assumed the percentage of patient visits for ILI is proportional to the number of infectives in the population. We estimated the associated parameters values so that the curve for the weekly number of infectives looks like the curve for the weekly percentage of patient visits for ILI shown in Fig. 1 (from North Carolina Public Health (2011)), which we refer to as the ILI% curve.

Because the vaccination coverage in North Carolina in the 2010–2011 flu season was about 45%, we set the vaccination coverage at 45% in the disease spread model throughout the calibration process. We used the following characteristics of the ILI% curve to calibrate the disease spread model (in Fig. 1, we let the week ending in 10/9/2010 be week 1):

1. Horizon of the epidemic: the flu season in North Carolina is October through May (North Carolina Public Health, 2012).
2. Peak of the epidemic: as suggested by the ILI% curve, the peak of the flu activity occurs at around week 19.
3. The epidemic take-off time: as suggested by the ILI% curve, the epidemic takes off at around week 11.
4. The time at which the epidemic activity returns to its initial level: as suggested by the ILI% curve, the epidemic activity returns to its initial level at around week 29.
5. The time at which the epidemic activity is minimum: as suggested by the ILI% curve, the epidemic activity is minimum at around week 33.
6. The ratio of the number of infectives at the peak to the initial number of infectives: as suggested by the ILI% curve, this ratio is about 10 to 1. It is also true for the 2007–2008 and 2008–2009 flu seasons (North Carolina Public Health, 2011).

Table 1
Model parameters and variables.

Parameter or variable	Description	Base case value of parameters
N	Number of regions (counties of North Carolina)	100
x_i	Number of vaccine doses allocated to region i in Phase-I	Decision Variable
y_i	Number of vaccine doses allocated to region i in Phase-II	Decision Variable
v_0	Minimum vaccination coverage in all regions in Phase-I	Decision Variable
ART	Attack rate threshold	Varies (5–15%)
$F_i(v_0)$	Probability of epidemic containment in region i with Phase-I vaccination coverage v_0	Depends on v_0 and *ART*
ξ_i	Bernoulli random variable indicating epidemic containment in region i as a result of Phase-I vaccination	Random variable
α	Vaccination coverage in previous flu seasons (in North Carolina)	45% (CDC, 2011c, 2010)
P_i	Population of region i	2010 US Census values
n_i	Lower bound on doses allocated to region i in Phase-I	$v_0 P_i$
m_i	Lower bound on doses allocated to region i in Phase-II	αP_i
r	Percentage increase in vaccination cost in Phase-II	Varies (0–60%)
c_i	Per-dose vaccination cost in region i in Phase-I	\$10 (CDC, 2012a)
d_i	Per-dose vaccination cost in region i in Phase-II	$(1 + r)c_i$
V_1	Number of vaccine doses made available in Phase-I	3,857,486
V_2	Number of vaccine doses made available in Phase-II	433,482

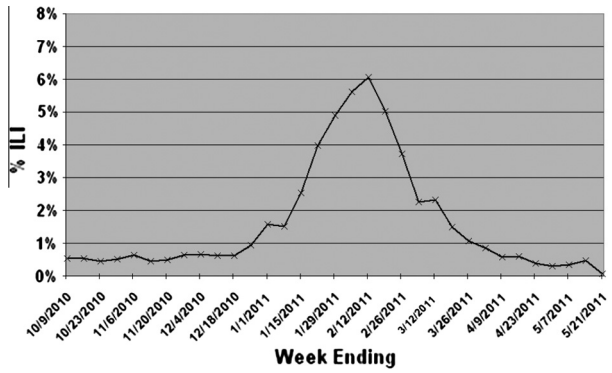

Fig. 1. Weekly percentage of patient visits for ILI in flu season 2010–2011 in North Carolina.

7. Rates of increase/decrease in ILI% around the peak: after the epidemic reaches its peak (at week 19) its activity decreases faster than it had increased before the peak (i.e., the absolute value of the slope of the ILI% curve is larger in the right side than the left side of the peak). This is also true for the 2007–2008 and 2008–2009 flu seasons (North Carolina Public Health, 2011). We assume this fast decrease in flu activity is due to the change in people's social behavior as the society copes with the epidemic (Larson, 2007) when the flu activity is maximum at week 19.

According to the ILI percentage data, flu activity is fairly low from October to mid-December. Flu activity increases in the subsequent weeks (which coincides with the winter holiday season) until it peaks in early February. This suggests a period starting in mid-December in which the exposure and infection rates increase. This increase could be due to importation of infectives, staying more indoors, etc. But from the modeling point of view, these have a similar effect, which is an increase in the exposure rate or horizontal incidence. We have assumed that this increase in the exposure rate is due to importation of infectives into different regions from outside (this could include importation from other neighboring regions). This importation of infectives may be due to travel to the given county from other counties or from outside of North Carolina. The annual person-trip volume for North Carolina (including internal visits) is about 36.8 million (Division of Tourism & Sports Development, 2011). We estimate the average daily visits from this volume and distribute it among different counties proportional to their population. We incorporate the infective importation rate from mid-December to the end of February. This rate is assumed to be highest in the middle of this period (i.e., mid-January), estimated to be one infective in every ten visitors, and decreases linearly until it reaches zero in mid-December and the end of February. The maximum infective importation rate in mid-January is 4.66 infectives per day in the Tyrrell county and 972.35 infectives per day in the Mecklenburg county.

Motivated by the ILI% curve, we also consider the change in the people's social behavior as society copes with the epidemic (Larson, 2007), especially at the peak of the flu activity. According to our estimate, the effective contact rate, and hence R_0, is reduced by 25% after the peak of the flu activity is reached in early February as a result of a higher level of precautions among individuals.

The initial number of infectives is estimated to be 4 infectives among every 10,000 individuals. Therefore there would be 2 initial infectives in Tyrrell county and 368 initial infectives in Mecklenburg county.

Fig. 2 presents the weekly number of infectives generated by the calibrated disease spread model in a single realization for a population of 1,000,000 (corresponding to Mecklenburg county with population 919,628) and 5000 (corresponding to Tyrrell county with population 4407). As expected, the stochastic effects are magnified in the smaller population resulting in an irregular curve in comparison with the smooth curve for the larger population. However, it can be easily verified that the curve for weekly number of infectives for both populations have all of the above-mentioned characteristics observed in the ILI% curve.

3.2. Stochastic programming model

In this subsection we formulate the two-phase vaccine allocation problem as a stochastic program. With the objective of minimizing the Phase-I vaccination cost plus the expected Phase-II vaccination cost and lower bounds on Phase-I and Phase-II vaccination coverage (implicitly imposed by a specific level of *ART* as well as the maximum demand for vaccine doses) the problem can be formulated as the following 2-SLP:

$$\min z(\boldsymbol{x}) = \sum_{i=1}^{N} c_i x_i + \mathcal{Q}(\boldsymbol{x}) \tag{1-a}$$

$$\sum_{i=1}^{N} x_i \leqslant V_1 \tag{1-b}$$

$$x_i \geqslant n_i, \quad i \in \mathcal{N} \tag{1-c}$$

where $\boldsymbol{x} = (x_i),\ i \in \mathcal{N}$, represents the vector of the first stage decision variables and $\mathcal{Q}(\boldsymbol{x}) = E_{\xi}[Q(\boldsymbol{x}, \xi)]$, where $\xi = (\xi_i),\ i \in \mathcal{N}$, represents the vector of N Bernoulli random variables $\xi_i,\ i \in \mathcal{N}$, with finite support $\Xi = \{(\delta_i), i \in \mathcal{N} | \delta_i = 0 \text{ or } 1\}$ and scenarios indexed by ω, and

$$Q(\boldsymbol{x}, \xi) = \min \sum_{i=1}^{N} d_i y_i(\omega) \tag{2-a}$$

$$\sum_{i=1}^{N} y_i(\omega) \leqslant V_1 + V_2 - \sum_{i=1}^{N} x_i \tag{2-b}$$

$$\xi_i x_i + y_i(\omega) \geqslant \xi_i m_i, \quad i \in \mathcal{N} \tag{2-c}$$

$$y_i(\omega) \geqslant 0, \quad i \in \mathcal{N} \tag{2-d}$$

where $\boldsymbol{y}(\omega) = (y_i(\omega)),\ i \in \mathcal{N}$, represents the vector of the second stage decision variables associated with scenario ω. Constraints (1-b) and (2-b) represent the limit on the vaccine doses available in Phase-I and Phase-II, respectively. Constraints (1-c) and (2-c) reflect the lower bound on the number of vaccine doses administered in region i in Phase-I and Phase-II, respectively. Constraint (2-c) is redundant if $\xi_i = 0$ (i.e., if the epidemic has been contained in Phase-I) and becomes $x_i + y_i(\omega) \geqslant m_i$ if $\xi_i = 1$ for scenario ω, which represents the lower bound $m_i - x_i$ on $y_i(\omega)$ to have a total of m_i vaccine doses administered in region i in Phase-II if the epidemic is not contained in Phase-I.

An important input of the developed 2-SLP model is the joint probability distribution of different scenarios. Due to the independence of $\xi_i,\ i \in \mathcal{N}$, the joint probability distribution of the scenarios in Ξ is the product of the Bernoulli distributions of $\xi_i,\ i \in \mathcal{N}$. Therefore we only need to estimate $F_i(v_0),\ i \in \mathcal{N}$, for each *ART* and v_0 level using the disease spread model.

There are $|\Xi| = 2^N$ different scenarios for the developed 2-SLP. Therefore, the number of scenarios grows exponentially with the number of regions, which might cause computational difficulties. We exploit the special problem structure to find the exact optimal Phase-I allocation very efficiently. In particular, we show that the optimal solution to the 2-SLP in (1) can be found by solving an LP with a similar size to that of the first stage problem. First, we

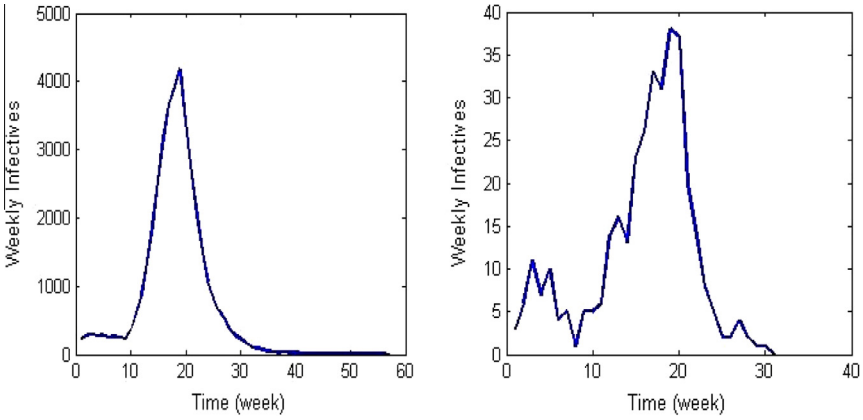

Fig. 2. Weekly number of infectives generated by the calibrated disease spread model for population of 1,000,000 (left) and 5000 (right).

establish a property of the optimal Phase-I allocation as follows. Let $\boldsymbol{x}^* = (x_i^*),\ i \in \mathcal{N}$, be an optimal Phase-I allocation for all scenarios or a subset of scenarios in Ξ.

Proposition 1. $x_i^* \leqslant m_i,\ i \in \mathcal{N}$.

Proof. If in an optimal allocation $\bar{\boldsymbol{x}} = (\bar{x}_i),\ i \in \mathcal{N}$, we have $\bar{x}_j > m_j$ for some $j \in \mathcal{N}$, then the solution obtained by letting $\bar{x}_j = m_j$ would be feasible with respect to constraint (2-c) and less costly with respect to the first stage objective function (1-a), contradicting the optimality of $\bar{\boldsymbol{x}}$. □

We use Proposition 1 to efficiently evaluate the function $Q(\boldsymbol{x}, \xi)$ at optimal for a known scenario ω as follows.

Proposition 2. *At an optimal Phase-I allocation $\boldsymbol{x}^*$, the solution to the recourse problem (2) for a given scenario ω is $y_i^* = \xi_i(m_i - x_i^*),\ i \in \mathcal{N}$, with the optimal objective function value $Q(\boldsymbol{x}, \boldsymbol{\xi}) = \sum_{i=1}^{N} \xi_i(m_i - x_i^*)d_i$.*

Proof. From (2-c) and Proposition 1 we have

$$y_i \geqslant \xi_i m_i - \xi_i x_i^*, \quad i \in \mathcal{N} \Rightarrow \begin{cases} y_i \geqslant m_i - x_i^* \geqslant 0, & \xi_i = 1 \\ y_i \geqslant 0, & \xi_i = 0 \end{cases} \tag{3}$$

If we ignore constraint (2-b) and note that $d_i \geqslant 0, i \in \mathcal{N}$, then the minimum value for the objective function (2-a) is obtained by setting $y_i^* = 0$ if $\xi_i = 0$ and $y_i^* = m_i - x_i^*$ if $\xi_i = 1$, which is the solution given by the proposition. Also $\boldsymbol{y}^* = (y_i^*), i \in \mathcal{N}$, given by the proposition satisfies constraint (2-b) because we have

$$\sum_{i=1}^{N} y_i^* \leqslant \sum_{i=1}^{N} (m_i - x_i^*) \leqslant V_1 + V_2 - \sum_{i=1}^{N} x_i^* \tag{4}$$

where the first inequality follows from Proposition 1 and the last inequality follows from the assumption $\sum_{i=1}^{N} m_i \leqslant V_1 + V_2$. Therefore $\boldsymbol{y}^* = (y_i^*), i \in \mathcal{N}$ given by the proposition is the optimal solution of the recourse problem (2) at an optimal Phase-I allocation x^*. □

Using Proposition 2 the recourse function at the optimal Phase-I allocation can be calculated directly.

Theorem 1. *At the optimal Phase-I allocation, the recourse function $\mathcal{Q}(x)$ can be calculated as follows,*

$$\mathcal{Q}(\boldsymbol{x}) = \sum_{i=1}^{N} \bar{\xi}_i (m_i - x_i) d_i \tag{5}$$

where $\bar{\xi}_i = E[\xi_i] = 1 - F_i(v_0),\ i \in \mathcal{N}$.

Proof. Using Proposition 2 we have

$$\mathcal{Q}(\boldsymbol{x}) = E_\xi[Q(\boldsymbol{x}, \boldsymbol{\xi})] = E_\xi\left[\sum_{i=1}^{N} \xi_i (m_i - x_i) d_i\right] = \sum_{i=1}^{N} E_\xi[\xi_i (m_i - x_i) d_i]$$
$$= \sum_{i=1}^{N} E_\xi[\xi_i](m_i - x_i) d_i = \sum_{i=1}^{N} \bar{\xi}_i (m_i - x_i) d_i. \quad \square$$

Using Theorem 1 we can solve the following LP instead of the 2-SLP in (1) to find the optimal Phase-I allocation.

$$\min z(\boldsymbol{x}) = \sum_{i=1}^{N} (c_i - \bar{\xi}_i d_i) x_i \tag{6-a}$$

$$\sum_{i=1}^{N} x_i \leqslant V_1 \tag{6-b}$$

$$x_i \geqslant n_i, \quad i \in \mathcal{N}. \tag{6-c}$$

Next, we provide insight into the optimal Phase-I allocation under a particular scenario. The fact that $d_i, i \in \mathcal{N}$, are nonincreasing in the region index i (see Assumption 1) can be used to efficiently find the optimal Phase-I allocation for a given scenario ω denoted by $\boldsymbol{x}^*(\omega) = (x_i^*(\omega)),\ i \in \mathcal{N}$. In general, the goal is to administer the required vaccine doses in Phase-I to avoid Phase-II vaccination, which is potentially more costly than Phase-I.

Theorem 2. *For a given scenario ω the optimal Phase-I allocation is*

$$x_i^*(\omega) = \begin{cases} n_i, & \xi_i = 0 \\ n_i + \min\{m_i - n_i, T_i\}, & \xi_i = 1 \end{cases} \tag{7}$$

where T_i represents the remaining additional vaccine doses for regions $i, i+1, \ldots, N$, i.e.,

$$T_i = A - \sum_{j=1}^{i-1} \xi_i (\min\{m_j - n_j, T_j\}), \quad i \in \mathcal{N} \tag{8}$$

where $A = V_1 - \sum_{i=1}^{N} n_i$.

Proof. First n_i vaccine doses are allocated to regions $i,\ i \in \mathcal{N}$, as the minimum requirement, leaving $A = V_1 - \sum_{i=1}^{N} n_i$ vaccine doses in Phase-I. If the epidemic is contained in region i in Phase-I (i.e., $\xi_i = 0$), no additional vaccine doses are required and $x_i^*(\boldsymbol{\omega}) = n_i$. Consider the regions in which the epidemic is not contained in Phase-I and denote the set of indices of these regions by J, that is $J = \{i \in \mathcal{N} | \xi_i = 1\}$. It follows from Proposition 1 that it is not

optimal to allocate more than m_i vaccine doses to regions $i, i \in J$. The vaccination cost is then minimized by prioritizing regions i, $i \in J$ for receiving additional $m_i - n_i$ vaccine doses over the minimum n_i vaccine doses in Phase-I in decreasing order of the second phase vaccination cost which results in the allocation given by (7) (note that $d_1 \geqslant d_2 \geqslant \cdots \geqslant d_N$ by Assumption 1). □

We use Theorem 2 in calculating *VSS* and *EVPI*, where we need to find the optimal Phase-I allocation for a reference scenario, and all scenarios, respectively. Using Theorem 2 decreases the computation time considerably as we do not need to solve an LP to find the optimal Phase-I allocation.

Remark. Although the allocation in Theorem 2 is optimal from the cost minimization point of view, it does not necessarily provide the most equitable allocation. Therefore in practical applications, the additional vaccine doses can be allocated to different regions proportional to their population. However, a different allocation will not change our numerical results for *VSS* and *EVPI* as $d_1 = d_2 = \cdots = d_N$ in our numerical experiments based on the available data.

3.3. Newsvendor model

In this subsection we formulate the two-phase vaccine allocation problem as a Newsvendor model with capacity constraints, which will allow us to give a closed-form solution for Phase-I allocation. We look at the vaccine allocation in each region as the demand for vaccine doses. The probability of epidemic containment in region i with Phase-I vaccination is assumed to remain at $F_i(v_0)$ even for a Phase-I coverage larger than v_0. This means that there is no incentive for Phase-I allocation, x_i, to be greater than the minimum required level, n_i, except for avoiding the possible (potentially more costly) Phase-II vaccination in case the epidemic is not contained with Phase-I vaccination (with probability $1 - F_i(v_0)$). But in the latter case, it is optimal to order all required doses, m_i, in Phase-I (note that according to Proposition 1 the optimal doses in Phase-I do not exceed m_i). Therefore, the demand for vaccine doses in Phase-I in region i is n_i with probability $F_i(v_0)$, and m_i with probability $1 - F_i(v_0)$.

Let $C_o^{(i)}$ and $C_u^{(i)}$ denote the unit cost of overstocking (i.e., ordering too many doses) and understocking (i.e., ordering too few doses) in Phase-I in region i. It is easy to verify that in our Newsvendor model

$$C_o^{(i)} = c_i \quad \text{and } C_u^{(i)} = d_i - c_i, \quad i \in \mathcal{N}. \tag{9}$$

Following the Newsvendor approach, we calculate the critical ratio, CR_i, as follows,

$$CR_i = \frac{C_u^{(i)}}{C_o^{(i)} + C_u^{(i)}} = \frac{d_i - c_i}{d_i} = 1 - \frac{c_i}{d_i} = \frac{r}{1+r}, \quad i \in \mathcal{N} \tag{10}$$

As a consequence of the two-point demand distribution, the unconstrained optimal Phase-I allocation will be $x_i = n_i$ if $F_i(v_0) \geqslant r/(1+r)$, and $x_i = m_i$ otherwise. To find the optimal Phase-I allocation with limited doses in Phase-I, we partition the set of regions, $\mathcal{N}$, into two subsets as follows,

$$\mathcal{N}^+ = \left\{ i \in \mathcal{N} | F_i(v_0) \geqslant \frac{r}{1+r} \right\} \quad \text{and } \mathcal{N}^- = \mathcal{N} \setminus \mathcal{N}^+ \tag{11}$$

Theorem 3. *The optimal Phase-I allocation is*

$$x_i^* = \begin{cases} n_i, & i \in \mathcal{N}^+ \\ n_i + \min\{m_i - n_i, T_{\kappa(i)}\}, & i \in \mathcal{N}^- \end{cases} \tag{12}$$

where $\kappa(i)$ represents the order of region i if the regions $i \in \mathcal{N}^-$ are reordered as $i_1, i_2, \ldots, i_{|\mathcal{N}^-|}$ such that

$$(1 - F_{i_1}(v_0))(d_{i_1} - c_{i_1}) \geqslant (1 - F_{i_2}(v_0))(d_{i_2} - c_{i_2}) \geqslant \cdots \geqslant (1 - F_{i_{|\mathcal{N}^-|}}(v_0))(d_{i_{|\mathcal{N}^-|}} - c_{i_{|\mathcal{N}^-|}}) \tag{13}$$

and $T_{\kappa(i)}$ represents the remaining additional vaccine doses for regions $i_{\kappa(i)}, i_{\kappa(i)+1}, \ldots, i_{|\mathcal{N}^-|}$, i.e.,

$$T_{\kappa(i)} = A - \sum_{j=1}^{\kappa(i)-1} \min\{m_{i_j} - n_{i_j}, T_{i_j}\}, \quad i \in \mathcal{N}^- \tag{14}$$

Proof. First n_i vaccine doses are allocated to regions $i, i \in \mathcal{N}$ as the minimum requirement, leaving $A = V_1 - \sum_{i=1}^{N} n_i$ vaccine doses in Phase-I. For $i \in \mathcal{N}^+$, the unconstrained optimal Phase-I allocation is $x_i = n_i$. Since n_i is the minimum required doses in Phase-I, the constrained optimal Phase-I allocation will also be $x_i^* = n_i$ for $i \in \mathcal{N}^+$. For $i \in \mathcal{N}^-$, the expected loss for each unit not ordered in Phase-I (i.e., the *risk*) is $(1 - F_i(v_0))(d_i - c_i)$, which represents the product of the probability that additional doses are required in Phase-II and the cost of understocking in Phase-I. Therefore the expected total loss, or equivalently the expected total cost, is minimized by prioritizing regions $i, i \in \mathcal{N}^-$, for receiving the additional $m_i - n_i$ vaccine doses over the minimum n_i vaccine doses in Phase-I in decreasing order of the expected loss represented in (13). □

Theorem 3 gives the optimal Phase-I allocation in a closed form, which is the advantage of the Newsvendor formulation. This alternative formulation provides additional insight into the optimal Phase-I allocation.

4. Results and discussion

We used MATLAB (version 7.8.0.347 R2009a) for our numerical experiments. All experiments were run on an Intel(R) Core(TM)2 Quad CPU with 8.00 gigabyte RAM. We estimated the probability of epidemic containment, $F_i(v_0), i \in \mathcal{N}$, for different values of v_0 (ranging from 0% to 45%) and *ART* (5%, 10%, and 15%) for each of the $N = 100$ counties of North Carolina using 5 batches with 200 simulation runs of the disease spread model in each batch. This took about 20 hours for all 100 counties. We use the sign ± to indicate the 95% confidence interval for the associated output. As an illustration, we have plotted $F_i(v_0)$ for *ART* = 10% for different values of v_0 for three counties including the smallest and the largest counties in Fig. 3. As Fig. 3 shows, the probability of epidemic containment heavily depends on the county population.

Assuming a vaccination cost of $10 per dose, the current continuous vaccination policy with a coverage of 45% of North Carolina population incurs a cost of $10 × ($V_1$ + V_2) = $42,909,680, or about $42.9 million. We report the cost savings in a similar way with only three significant digits. The expected total cost for different values of v_0, *ART*, and r are presented in Fig. 4. In the next two subsections, we discuss the impact of different values of v_0, *ART*, and r on the vaccine allocation and the expected total cost.

4.1. Impact of v_0 and ART

When v_0 is small, say $v_0 \leqslant 5\%$, the probability of epidemic containment with Phase-I vaccination for all values of *ART* is small in all counties. Therefore all V_1 available vaccine doses in Phase-I and, on average, a considerable amount of additional V_2 vaccine doses available in Phase-II are allocated to different regions yielding a statewide coverage (i.e., the ratio of expected vaccine doses allocated in Phase-I and Phase-II to the state population) of 44.5%.

For *ART* = 5%, the probability of epidemic containment for $v_0 \leqslant 30\%$ remains small in different counties. Therefore we

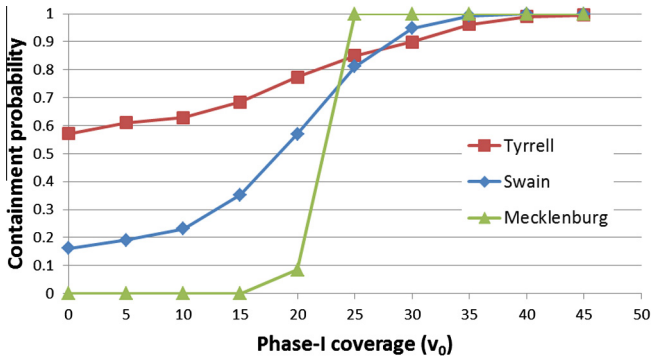

Fig. 3. $F_i(v_i)$ for $ART = 10\%$ for Tyrrell, Swain, and Mecklenburg counties with populations 4407, 13,981, and 919,628, respectively.

observe a high statewide coverage of 44.6% for $v_0 \leqslant 30\%$. However, for larger values of v_0 the probability of epidemic containment is large for $ART = 5\%$, 10%, or 15% in different counties. For example, for $v_0 = 40\%$, the probability of epidemic containment for different values of ART is at least 0.717 ± 0.041 in all counties. Thus, in Phase-I the minimum required coverage of $v_0 = 40\%$ is allocated to all counties with a small chance of Phase-II allocation yielding a statewide coverage of 40%.

We observe that $v_0 = 40\%$ is indeed the optimal Phase-I coverage for $ART = 5\%$ as demonstrated by Fig. 4. The optimal Phase-I coverage for different values of ART are presented in Table 2. We have also reported the statewide coverage. As Table 2 shows, with $v_0^* = 40\%$ about $0.05 \times 9{,}535{,}483 = 476{,}774$ vaccine doses could be saved by implementing the two-phase vaccination policy with an associated monetary saving of about $4 million. Note that the case of $v_0^* = 40\%$, which is the maximum possible coverage in Phase-I, shows the importance of early mass vaccination before October in comparison with the continuous vaccination policy currently implemented. In fact our results show that with 40% coverage before the beginning of the flu season in October, the attack rate will be below 10% in all counties with a probability of at least 0.989 ± 0.005. Furthermore, the statewide attack rate (i.e., the percentage of the state population who become infective during the flu season) will be 4.44%. Therefore the additional vaccine doses for 5% of the population with the corresponding cost of $4.8 million could be saved by early vaccination.

The curves of expected total cost for $ART = 10\%$ and 15% are similar. Both have a linear portion at the beginning corresponding to small values of v_0 where the probability of epidemic containment is small and therefore almost all available vaccine doses are allocated in Phase-I and Phase-II yielding a large cost close to the cost of the continuous vaccination policy. After the initial linear portion, the expected total cost decreases for larger values of v_0 until it

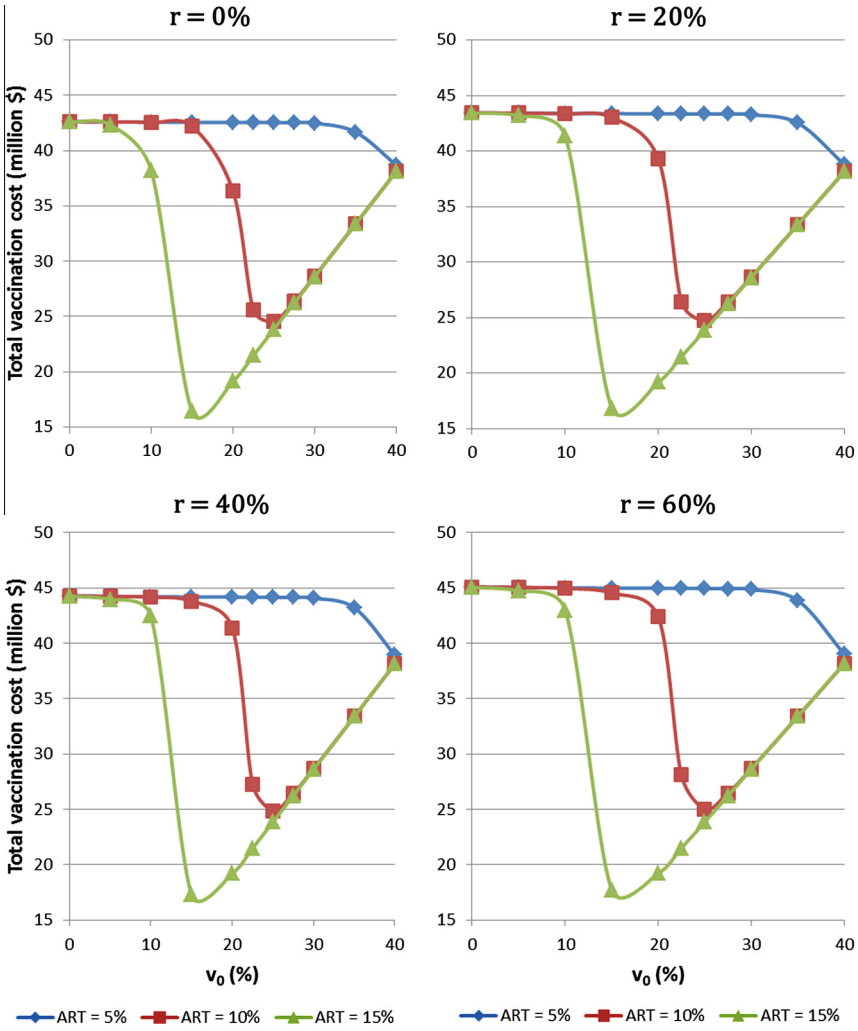

Fig. 4. Expected total cost for different values of v_0 (minimum Phase-I coverage), ART (attack rate threshold), and r (percentage increase in vaccination cost in Phase-II).

Table 2
Optimal vaccination coverage.

ART (%)	Optimal phase-I coverage (v_0^*) (%)	Expected doses saved	Expected monetary saving ($ million)[a]	Statewide coverage (%)	Statewide attack rate (%)
5	40	476,774	[4.2,3.9]	40.6	4.44
10	25	1,830,813	[18.3,17.9]	25.8	8.27
15	15	2,650,864	[26.5,25.2]	17.2	13.04

[a] The reported values are for $r = 0\%$ and $r = 60\%$, respectively.

reaches its minimum as reported in Table 2. In this region of the expected total cost curve, as v_0 increases, the probability of epidemic containment in Phase-I increases. Therefore, on average, fewer vaccine doses are required to contain the epidemic, hence the decrease in the expected total cost.

As Table 2 demonstrates, with optimal Phase-I coverage a considerable number of vaccine doses could be saved with an associated large monetary savings and an acceptable attack rate. The optimal values for v_0 for different values of *ART* have a characteristic which makes these allocations more plausible for health care officials from an equity point of view: under these allocations the Phase-I coverage in all counties will be the same, and all equal to v_0^*.

After the minimum cost point, the expected total cost for *ART* = 10% and 15% increases linearly with v_0. The reason is that at the minimum point, the probability of epidemic containment with Phase-I coverage is sufficiently large that no additional vaccination in Phase-I is necessary over the minimum required coverage of v_0^*. Therefore increasing v_0 over v_0^* merely results in unnecessary vaccination in Phase-I incurring the associated excess cost. As expected, the final linear parts of the curves of the expected total cost for *ART* = 10% and 15% overlap as the associated probabilities of epidemic containment with Phase-I vaccination are both very close to 1.

Remark. The actual monetary savings are larger than our estimates because we have not considered other relative costs such as work loss, treatment, and hospitalization. Indeed, these costs and also the number of deaths depend on the attack rate. Therefore by limiting the attack rate, we have actually limited these other costs as well. This is based on the purpose of our research, which is to demonstrate that the epidemic can be effectively "contained" with fewer vaccine doses if the outcome of vaccination is evaluated. Therefore the target is to contain the epidemic (i.e., keep the attack rate below *ART*, which also limits other costs and the number of deaths) with the minimum number of vaccine doses.

4.2. Impact of $\boldsymbol{r}$

As Fig. 4 demonstrates, the expected total cost for each value of *ART* has the same pattern for different values of *r*. We have only considered $r \geqslant 0\%$ cases in our numerical experiments because these cases are of greater interest (as discussed previously, the optimal allocation for $r < 0\%$ is trivial). In fact for $r < 0$ the optimal Phase-I allocation and the expected allocation in Phase-II, and hence the expected total doses allocated in both phases, the statewide coverage, and the expected attack rate will all be exactly the same as for $r = 0\%$. The only difference is the expected total cost which can be easily calculated based on the value of *r* and the expected Phase-II allocation.

The optimal Phase-I allocation (reported in Table 2) for different values of *ART* does not change when *r* is increased from 0% to 60%. This may in part be due to the small change in the expected total cost for different values of *r*. Therefore the optimal Phase-I allocation is robust to changes in *r*, at least in the considered range of 0–60%.

The expected total allocated doses (equivalently the statewide coverage) and the expected total cost are nondecreasing in *r* for each fixed value of v_0 and *ART*. The reason is that as *r* increases, the optimal number of vaccine doses allocated in Phase-I to a particular region either remains unchanged or increases in order to administer fewer vaccine doses in Phase-II which now has become more costly with an increase in *r*. Therefore the associated total vaccination cost either remains unchanged or increases.

Remark. Currently we have estimated V_1 based on the vaccine share of North Carolina in the early months of the flu season (i.e., August through October), which is about 40% of the state population, to be able to compare the two-phase policy with the continuous policy. But as our numerical results demonstrate, with *ART* = 10% or 15% considerably fewer vaccine doses are required in Phase-I. Therefore our model and the numerical results can indeed be used to estimate the optimal value of $V_1$and V_2. Based on the value of *ART*, the optimal value for V_1 can be estimated as $V_1^* = v_0^* P$, where *P* represents the state population. The optimal value of V_2 can be estimated as the difference between Phase-I coverage and the statewide coverage. For example, for *ART* = 15% an estimate for the optimal value of V_1 and V_2 is $V_1^* = (15\%)P$ and $V_2^* = (17.2\% - 15\%)P = (2.2\%)P$. Therefore in practice, $V_1^* + V_2^*$ doses can be ordered for production prior to the flu season. Note that due to the aggregation of demand in different regions, the variability associated with calculating the optimal order quantity is considerably reduced.

After Phase-I, health care officials decide whether to allocate more doses to each region based on the *observation of flu activity in each region*. In other words, the 2-SLP model (similar to any other two-stage stochastic problem) or the Newsvendor model are used to find the optimal allocation in Phase-I under uncertainty. For Phase-II, the uncertainty has been revealed and the Phase-II allocation is determined accordingly.

4.3. VSS and EVPI

In this subsection we calculate *VSS* and *EVPI* using the 2-SLP model. To evaluate the performance of some typical allocations, we calculate *VSS* as the percentage reduction in the minimum cost if the optimal allocation x^* (obtained from Theorem 3 or by solving problem (6)) is used compared to the optimal allocation corresponding to a specific reference scenario. We consider three reference scenarios and calculate the expected value of the reference scenario (*EVRS*). We consider the worst case scenario $\boldsymbol{\xi}^w(\xi_i^w = 1, i \in \mathcal{N})$, the best case scenario $\boldsymbol{\xi}^b(\xi_i^b = 0, i \in \mathcal{N})$, and the round scenario $\boldsymbol{\xi}^r(\xi_i^r = round(1 - F_i(v_0))$, where $round(a)$ returns the closest integer to *a*. Note that in the round scenario it is assumed that the epidemic is contained in Phase-I only in regions with epidemic containment probability greater than 0.5. For each of $\boldsymbol{\xi}^w, \boldsymbol{\xi}^b$, and $\boldsymbol{\xi}^r$, we use Theorem 2 to find the associated optimal solution denoted by $\boldsymbol{x}' = (x_i'), i \in \mathcal{N}$. Then we use Theorem 1 to calculate $\mathcal{Q}(\boldsymbol{x}') = E_{\xi}[Q(\boldsymbol{x}', \boldsymbol{\xi})]$ and the expected cost of solution $\boldsymbol{x}'$, that is,

$$EVRS = z(\boldsymbol{x}') = \sum_{i=1}^{N} (c_i - \bar{\xi}_i d_i) x_i' \tag{15}$$

Then we calculate *VSS* as follows,

$$VSS = 100 \times \frac{z(\boldsymbol{x}') - z(\boldsymbol{x}^*)}{z(\boldsymbol{x}')} \tag{16}$$

We denote the associated *VSS* of the worst case, best case, and round scenarios by VSS^w, VSS^b, and VSS^r, respectively, with the associated optimal allocations $\boldsymbol{x}^w$, $\boldsymbol{x}^b$, and $\boldsymbol{x}^r$.

Another useful criterion is *EVPI* which provides an estimate of the benefits of a surveillance system that forecasts the epidemic progression. We calculate *EVPI* as the percentage reduction in the minimum cost if exact information about the future was available. First, we calculate the wait and see solution (*WS*) as follows. We find the optimal solution for each scenario ω, denoted by $\bar{\boldsymbol{x}}(\omega)$, using Theorem 2. Then $WS = E_\omega[z(\bar{\boldsymbol{x}}(\omega))]$ and

$$EVPI = 100 \times \frac{z(\boldsymbol{x}^*) - E_\omega[z(\bar{\boldsymbol{x}}(\omega))]}{z(x^*)} \tag{17}$$

Due to computational limitations in calculating *WS*, in this subsection we only consider a subset of the North Carolina counties including ten of the counties with various population sizes (logarithmically evenly distributed) including the smallest and largest counties. These counties account for about 21.24% of the North Carolina population. Therefore we reduce the available vaccine doses in Phase-I and Phase-II proportionally for this experiment and set $V_1 = 819{,}443$ and $V_2 = 92{,}084$. We calculate VSS^w, VSS^b, VSS^r, and *EVPI* for different values of r in two cases: $v_0 = 25\%$ and $ART = 10\%$, and $v_0 = 20\%$ and $ART = 10\%$. We have chosen these cases based on the average attack rate in the United States, which is about 10% (Homeland Security, 2011; World Health Organization, 2009), and the optimal Phase-I coverage of $v_0^* = 25\%$ corresponding to $ART = 10\%$ (see Table 2). We refer to these cases as Case 1 and Case 2, respectively. In fact the probability of epidemic containment in each county for each of these cases has a characteristic which makes these cases of particular interest to us. The results are presented in Table 3. Using Theorem 2, it takes only about 5 seconds, on average, to calculate any of VSS^w, VSS^b, VSS^r, and *EVPI*.

First note that $VSS^b = 0\%$ for $r = 0\%$ in both cases. This implies that the optimal allocation for the best case scenario ξ^b (i.e., $x_i^b = n_i, i \in \mathcal{N}$) is indeed the optimal allocation for $r = 0\%$ which is consistent with our previous note in Section 3 regarding the optimal allocation in Phase-I if $c_i \geqslant d_i$, $i \in \mathcal{N}$ (or equivalently if $r \leqslant 0\%$). Furthermore, $\boldsymbol{x}^b$ is the optimal allocation for $r = 0\%$ even with perfect information about the Phase-I vaccination outcome. Therefore *EVPI* is equal to zero for $r = 0\%$ in both cases. Also *EVPI* is nondecreasing in r in both cases as a consequence of the increase in the Phase-II vaccination cost, that is, the increase in the uncertain future cost.

In Case 1, the probability of epidemic containment is so large in all counties ranging from 0.81 to 1.00 that it is optimal to allocate only the minimum required doses in Phase-I for all values of r. But this is indeed the optimal allocation for $\xi^b = \xi^r$, hence the optimality of $\boldsymbol{x}^b = \boldsymbol{x}^r$. As a result, we observe that $VSS^b = VSS^r = 0.00\%$ for all values of r. On the other hand, $\boldsymbol{x}^w$ performs poorly in Case 1 because it allocates all available vaccine doses in Phase-I while many counties might not really need the additional 20% coverage over the minimum required coverage of $v_0 = 25\%$ in Phase-I due to the large probability of epidemic containment in all counties. We observe that in Case 1 *EVPI* is very small for all values of r. This is due to the fact that the probability of epidemic containment in all counties is so large that we are *almost* certain about the future. In particular, the probability of epidemic containment in the four most populated counties, which accommodate about 90% of the population of the considered ten counties, is almost 1. Therefore perfect information in this case is not very useful.

In Case 2, the probability of epidemic containment in the ten considered counties in the increasing order of population are 0.774, 0.629, 0.570, 0.485, 0.405, 0.338, 0.291, 0.243, 0.151, and 0.084. We observe that the performance of $\boldsymbol{x}^b$ becomes worse as r increases because it is assumed that the epidemic will be contained in all counties while with a high probability it will not be in at least the four most populated counties and therefore a large number of doses will be required in Phase-II in those counties. The performance of $\boldsymbol{x}^r$ and $\boldsymbol{x}^w$ improves as r increases from 0% to 20% to 40% because with a higher Phase-II vaccination cost it is more reasonable to allocate all available vaccine doses in Phase-I which is the case for $\boldsymbol{x}^r$ and $\boldsymbol{x}^w$. However, the performance of $\boldsymbol{x}^w$ and $\boldsymbol{x}^r$ becomes worse as r increases from 40% to 60%. This is due to the priority of smaller counties, which have smaller indices in our numerical experiment, to receive additional vaccine doses in Phase-I over the minimum required doses while the larger counties will probably need additional doses in Phase-II.

Table 3
Value of stochastic solution (*VSS*) and expected value of perfect information (*EVPI*) for selected cases.

		$r = 0\%$	$r = 20\%$	$r = 40\%$	$r = 60\%$
Case 1:	VSS^b	0.00	0.00	0.00	0.00
$v_0 = 25\%$	VSS^r	0.00	0.00	0.00	0.00
$ART = 10\%$	VSS^w	37.63	37.52	37.41	37.29
	EVPI	0.00	0.18	0.36	0.54
Case 2:	VSS^b	0.00	2.76	9.22	15.20
$v_0 = 20\%$	VSS^r	8.21	3.37	3.15	3.54
$ART = 10\%$	VSS^w	8.58	3.84	3.68	4.14
	EVPI	0.00	5.84	6.86	7.27

In Case 2 we can say $\boldsymbol{x}^b$ has the best performance for smaller values of r (i.e., 0% and 20%) while $\boldsymbol{x}^r$ has the best performance for larger values of r (i.e., 40% and 60%). This is independent of the probabilities of epidemic containment in different regions. We observe that in Case 2 *EVPI* has a considerably higher value than Case 1 for all values of r except for $r = 0\%$. The reason is that in contrast to Case 1, in Case 2 the probability of epidemic containment in none of the ten considered counties is close to 1. Therefore, identifying which counties will require Phase-II vaccination is valuable information.

In general, information about epidemic containment in larger counties is more valuable than smaller counties as larger counties require more vaccine doses. As supported by our numerical experiments, when the Phase-I coverage is sufficiently high (more than 25% according to our results) that the probability of epidemic containment, particularly in larger counties, is close to 1, *EVPI* is small. However, when the Phase-I coverage is relatively low, information about the Phase-I vaccination outcome becomes more valuable and could reduce the expected total cost considerably. Therefore, depending on the Phase-I coverage, epidemic alerting and detection systems could be valuable aids for planning vaccine allocation.

4.4. A Heuristic for vaccine allocation

Based on the values of *VSS* for the worst case, best case, and round scenarios, which are measures for the performance of the associated allocations, we propose a heuristic for vaccine allocation in Phase-I. If the increase in vaccination cost in Phase-II is small (less than 20% according to our numerical results), a good heuristic solution for Phase-I allocation is $\boldsymbol{x}^b$, i.e., to allocate only the minimum required doses to all regions. However, if vaccination in Phase-II is considerably more costly than Phase-I (more than 20% according to our numerical results), a good heuristic solution for Phase-I allocation is $\boldsymbol{x}^r$, i.e., to allocate the minimum required doses to all regions and then to allocate the remaining doses to regions with epidemic containment probability less than 0.5. The additional vaccine doses allocated to these regions should be equal to the difference between the minimum requirements in Phase-I and Phase-II. To minimize the expected total cost, regions with a higher second phase vaccination cost should receive priority for

receiving the remaining doses. Our results show that the performance of this easy-to-implement Phase-I allocation is insensitive to the increase in vaccination cost in Phase-II provided this increase is not too small (more than 20% according to our numerical results). Allocation of all available vaccine doses in Phase-I (i.e., $\boldsymbol{x}^w$), which corresponds to the worst case scenario, is very close to the continuous vaccination policy currently implemented in the sense that under both policies all available vaccine doses are allocated without any evaluation of the vaccination outcome. Therefore, as our results show, this policy performs poorly if the Phase-I coverage is relatively high (e.g., in Case 1 of Table 3).

5. Conclusions

We proposed a two-phase vaccination policy to contain an epidemic with Phase-I in the early stages of the epidemic and Phase-II in the middle of the epidemic after observing the outcome of Phase-I vaccination. This two-phase vaccination policy not only results in a considerable reduction in the number of vaccine doses required to contain the epidemic (up to 50%), and hence the vaccine production cost, but also results in a lower attack rate and also allows for redistribution of vaccine doses after Phase-I to use the vaccine doses more efficiently.

We considered two alternative formulations of the vaccine allocation problem. We formulated the problem of optimal allocation of vaccine doses to different regions in each phase as a 2-SLP and used the special problem structure to reduce it to an LP with a similar size to that of the first stage problem. We obtained structural properties which helped speed up computation. We also presented a Newsvendor formulation of the vaccine allocation problem which allowed us to give a closed-form solution for the optimal allocation. The Newsvendor formulation provided additional insight into the optimal allocation.

To illustrate the advantage of the two-phase vaccination policy over the continuous vaccination policy we considered 100 counties of North Carolina in case of seasonal influenza. We solved the problem for different values of v_0 (Phase-I coverage), *ART* (attack rate threshold), and r (percentage increase in vaccination cost in Phase-II) and found the optimal Phase-I allocation for different values of *ART* and r. We developed a disease spread simulation model based on the stochastic SEIR model to estimate the probability of epidemic containment with Phase-I vaccination in each region for different values of v_0 and *ART*.

The attack rate and the expected cost depend on the value of *ART* determined by health care officials. For *ART* = 5%, 10%, and 15%, the corresponding optimal Phase-I allocations are vaccination of $v_0^* = 40\%$, 25%, and 15% of each county population, respectively. These vaccination policies yield a statewide attack rate of 4.44%, 8.27%, and 13.04%, while saving a total of 476,774, 1,830,813, and 2,650,864 vaccine doses, on average, and 3.8, 17.9, and 25.2 million dollars in vaccine production and administration, respectively. This is in comparison with the current continuous vaccination policy with a statewide coverage of 45% in North Carolina. The (expected) statewide coverage in each case (i.e., 40.6%, 25.8%, and 17.2%, respectively) can be used as an estimate for the optimal number of vaccine doses which should be ordered for production prior to the flu season. The abovementioned optimal Phase-I allocations for different values of *ART* assign the same coverage in Phase-I to all counties, and therefore, equity among different counties is respected. Further, our results show that the optimal Phase-I allocation is insensitive to the increase in vaccination cost in Phase-II in comparison with Phase-I at least in the considered range of r = 0–60%.

We used the developed 2-SLP to calculate *VSS* and *EVPI*. To calculate *VSS*, we considered three reference scenarios: the best case scenario (the epidemic is contained in Phase-I in all regions), the worst case scenario (the epidemic is not contained in Phase-I in any region), and the round scenario (the epidemic is contained in Phase-I only in regions with epidemic containment probability greater than 0.5). We used *VSS* to examine the performance of the associated optimal Phase-I allocations, and accordingly, to propose an effective heuristic for vaccine allocation.

In our model, *EVPI* gives an estimate of the value of accurate information about the outcome of Phase-I vaccination with respect to epidemic containment. It provides an upper bound on the value of epidemic detection and alerting systems that could help public health officials in determining how much should be invested for providing such systems. Our results show that epidemic alerting and detection systems are more valuable (i.e., the associated *EVPI* will be higher) if the vaccination coverage in Phase-I is relatively low.

When there is an epidemic in the community, usually other interventions, such as antiviral prophylaxis and treatment, and also non-pharmaceutical interventions, such as isolation, are implemented in addition to vaccination. In our model, the impact of these interventions are reflected in the value of the basic reproduction number, R_0, which determines the effective contact rate among individuals. Further, the cost of these interventions are related to the number of infectives, and hence, the attack rate. There might also be constraints on resources other than vaccine doses, such as human and facility resources, which are related to the attack rate as well. Therefore, *ART* can be used to account for these other constraints as well. Furthermore, *ART* can be used to limit the attack rate, which is of course an important decision factor in addition to cost.

There are some limitations to our study. First, we only considered two vaccination phases in our model. In reality, there could be more than two vaccination phases, in which case the problem can be formulated as a multi-stage stochastic program. Intuitively, with more vaccination phases the saving in vaccine doses and the associated monetary value should be larger provided that there is enough time between different phases to evaluate the impact of previous phases with sufficient accuracy. Nevertheless, our results provide conservative estimates of the potential benefits of using a sequential decision making process for vaccine allocation. A second limitation of our modeling approach is that in calculating the monetary benefit of early vaccination in Phase-I we only considered the monetary saving in vaccine production. However, there are other monetary benefits as a result of a lower attack rate, such as work loss, treatment, and hospitalization costs associated with the averted cases. This results in an underestimation of monetary benefits of the proposed two-phase vaccination policy. A third limitation is the way we have accounted for transportation of infectives between regions in the disease spread model. We did not explicitly simulate the individual-by-individual travel of infectives between regions. Instead, we considered a general daily travel rate into different regions. Finally, we assumed the probability of epidemic containment corresponding to the minimum required vaccination in Phase-I (i.e., $F_i(v_0),\ i \in \mathcal{N}$) for tractability. But the coverage in Phase-I might be higher than the minimum required level (i.e., v_0). However, this assumption is not restrictive with respect to finding the optimal allocation because at the optimal Phase-I coverage (i.e., v_0^*) only the minimum required doses are allocated to all regions. Therefore the probability of epidemic containment is indeed $F_i(v_0^*),\ i \in \mathcal{N}$, at the optimal coverage. This is also true for coverage levels higher than the optimal coverage.

Acknowledgments

We are grateful to the anonymous referees for their helpful and constructive comments. In particular, we are grateful to the referee

who brought the Newsvendor model to our attention. This research was carried out by the North Carolina Preparedness and Emergency Response Research Center (NCPERRC) which is part of the UNC Center for Public Health Preparedness at the University of North Carolina at Chapel Hill's Gillings School of Global Public Health and was supported by the Centers for Disease Control and Prevention (CDC) Grant 1PO1 TP 000296.

References

Ball, F. G., & Lyne, O. D. (2002). Optimal vaccination policies for stochastic epidemics among a population of households. *Mathematical Biosciences, 177–178*, 333–354.

Ball, F., Mollison, D., & Scalia-Tomba, G. (1997). Epidemics with two levels of mixing. *Annals of Applied Probability, 7*, 46–89.

Bauch, C. T., & Earn, D. J. (2004). Vaccination and the theory of games. *Proceedings of the National Academy of Sciences of the United States of America, 101*, 13391–13394.

Becker, N. G., & Starczak, D. N. (1997). Optimal vaccination strategies for a community of households. *Mathematical Biosciences, 139*, 117–132.

Brito, D. L., Sheshinski, E., & Intriligator, M. D. (1991). Externalities and compulsory vaccinations. *Journal of Public Economics, 45*, 69–90.

Centers for Disease Control and Prevention (2010). *Interim results: state-specific seasonal influenza vaccination coverage – United States. August 2009–January 2010. Morbidity and mortality weekly report* (Vol. 59. pp. 477–484).

Centers for Disease Control and Prevention (2011a). *Final state-level influenza vaccination coverage estimates for the 2010–2011 season–United States. National Immunization Survey and Behavioral Risk Factor Surveillance System. August 2010 through May 2011*. <http://www.cdc.gov/flu/professionals/vaccination/coverage_1011estimates.htm#1>. (Accessed on 07.03.12).

Centers for Disease Control and Prevention (2011b). *Flu vaccine effectiveness: questions and answers for health professionals*. <http://www.cdc.gov/flu/professionals/vaccination/effectivenessqa.htm> (Accessed on 05.12.11).

Centers for Disease Control and Prevention (2011c). *Interim results: state-specific influenza vaccination coverage — United States. August 2010–February 2011. Morbidity and mortality weekly Report*. (Vol. 60, pp. 737–743).

Centers for Disease Control and Prevention (2011d). *Seasonal influenza (Flu)*. <http://www.cdc.gov/flu/professionals/diagnosis> (Accessed on 05.12.11).

Centers for Disease Control and Prevention (2011e). *Seasonal influenza: Questions & answers*. <http://www.cdc.gov/flu/about/qa/disease.htm#us-flu-season> (Accessed on 07.03.12).

Centers for Disease Control and Prevention (2012a). *CDC vaccine price list*. <http://www.cdc.gov/vaccines/programs/vfc/cdc-vac-price-list.htm>. (Accessed on 09.03.12).

Centers for Disease Control and Prevention (2012b). *Seasonal influenza vaccine & total doses distributed*. <http://www.cdc.gov/flu/professionals/vaccination/vaccinesupply.htm> (Accessed on 08.03.12).

Chapman, G. B., & Coups, E. J. (1999). Predictors of influenza vaccine acceptance among healthy adults. *Preventive Medicine, 29*, 249–262.

Chick, S. E., Mamani, H., & Simchi-Levi, D. (2008). Supply chain coordination and influenza vaccination. *Operations Research, 56*, 1493–1506.

Cho, S. H. (2010). The optimal composition of influenza vaccines subject to random production yields. *Manufacturing & Service Operations Management, 12*, 256–277.

Chowell, G., Miller, M. A., & Viboud, C. (2008). Seasonal influenza in the United States, France, and Australia: transmission and prospects for control. *Epidemiology and Infection, 136*, 852–864.

Chowell, G., Nishiura, H., & Bettencourt, L. M. A. (2007). Comparative estimation of the reproduction number for pandemic influenza from daily case notification data. *Journal of the Royal Society Interface, 4*, 155–166.

Deo, S., & Corbett, C. J. (2008). Cournot competition under yield uncertainty: The case of the US influenza vaccine market. *Manufacturing & Service Operations Management, 11*, 563–576.

Division of Tourism, Film and Sports Development (2011). *2010 North Carolina Annual Report*. <http://www.annualreport.visitnc.com/year-in-review/2009-results/>. (Accessed on 14.03.12).

Flahault, A., Vergu, E., Coudeville, L., & Grais, R. F. (2006). Strategies for containing a global influenza pandemic. *Vaccine, 24*, 6751–6755.

Galvani, A. P., Reluga, T. C., & Chapman, G. B. (2007). Long-standing influenza vaccination policy is in accord with individual self-interest but not with the utilitarian optimum. *Proceedings of the National Academy of Sciences of the United States of America, 104*, 5692–5697.

Geoffard, P. Y., & Philipson, T. (1996). Rational epidemics and their public control. *International Economic Review, 37*, 603–624.

Gerdil, C. (2003). The annual production cycle for influenza vaccine. *Vaccine, 21*, 1776–1779.

Hethcote, H. W. (2000). The mathematics of infectious diseases. *SIAM Review, 42*, 599–653.

Hill, A. N., & Longini, I. M. Jr, (2003). The critical vaccination fraction for heterogeneous epidemic models. *Mathematical Biosciences, 181*, 85–106.

Homeland Security (2011). *Flu pandemic morbidity/mortality*. <http://www.globalsecurity.org/security/ops/hsc-scen-3_flu-pandemic-deaths.htm>. (Accessed on 08.03.12).

Khazeni, N., Hutton, D. W., Garber, A. M., Hupert, N., & Owens, D. K. (2009). Effectiveness and cost-effectiveness of vaccination against pandemic influenza (H1N1) 2009. *Annals of Internal Medicine, 151*, 829–839.

Kornish, L. J., & Keeney, R. L. (2008). Repeated commit-or-defer decisions with a deadline: The influenza vaccine composition. *Operations Research, 56*, 527–541.

Larson, R. (2007). Simple models of influenza progression within a heterogeneous population. *Operations Research, 55*, 399–412.

Larson, E. B., Olsen, E., Cole, W., & Shortell, S. (1979). The relationship of health benefits and a postcard reminder to influenza vaccination. *Journal of Family Medicine, 8*, 1207–1211.

Li, M. Y., Graef, J. R., Wand, L., & Karsai, J. (1999). Global dynamics of a SEIR model with varying total population size. *Mathematical Biosciences, 160*, 191–215.

North Carolina Public Health (2011). *North Carolina weekly influenza surveillance summary: 2010–2011 influenza season*, 33.

North Carolina Public Health (2012). *Diseases & topics: Influenza*. <http://www.epi.publichealth.nc.gov/cd/diseases/flu.html>. (Accessed on 14.03.12).

Philipson, T. (2000). Economic epidemiology and infectious diseases. *Handbook of Health Economics, 1*, 1761–1799.

Reluga, T. C., Bauch, C. T., & Galvani, A. P. (2006). Evolving public perceptions and stability in vaccine uptake. *Mathematical Biosciences, 204*, 185–198.

Tanner, M. W., Sattenspiel, L., & Ntaimo, L. (2008). Finding optimal vaccination strategies under parameter uncertainty using stochastic programming. *Mathematical Biosciences, 215*, 144–151.

U.S. Census Bureau (2012). *State & County QuickFacts*. <http://www.quickfacts.census.gov/qfd/states/37000.html>. (Accessed on 01.29.12).

World Health Organization (2009). *Assessing the severity of an influenza pandemic*. <http://www.who.int/csr/disease/swineflu/assess/disease_swineflu_assess_20090511/en/index.html>. (Accessed on 08.03.12).

Wu, J. T., Wein, L. M., & Perelson, A. S. (2005). Optimization of influenza vaccine selection. *Operations Research, 53*, 456–476.

Yarmand, H. (2010). *Cost-effectiveness analysis of different interventions for H1N1: What is the optimal level of vaccination and self-isolation in case of an H1N1 outbreak*. Saarbrücken, Germany: LAP LAMBERT Academic Publishing.

Yarmand, H., & Ivy, J. S. (in press). Optimal intervention strategies for stochastic susceptible-infective disease spread model: A household view. *Simulation: Transactions of the Society for Modeling and Simulation International*.

Yarmand, H., Ivy, J. S., & Roberts, S. D. (in press). Identifying optimal mitigation strategies for responding to a mild influenza epidemic. *Simulation: Transactions of the Society for Modeling and Simulation International*.

Yarmand, H., & Ivy, J. S. (2013). Analytic solution of the susceptible-infective disease spread model with state-dependent contact rates and different intervention strategies. *Simulation: Transaction of the Society for Modeling and Simulation International, 89*, 703–721.