Service Science

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

Engineering Effective Responses to Influenza Outbreaks

Stan N. Finkelstein, Richard C. Larson, Karima Nigmatulina, Anna Teytelman

To cite this article:
Stan N. Finkelstein, Richard C. Larson, Karima Nigmatulina, Anna Teytelman (2015) Engineering Effective Responses to Influenza Outbreaks. Service Science 7(2):119-131. http://dx.doi.org/10.1287/serv.2015.0099

SERVICE SCIENCE
Vol. 7, No. 2, June 2015, pp. 119–131
ISSN 2164-3962 (print) | ISSN 2164-3970 (online)

informs

http://dx.doi.org/10.1287/serv.2015.0099

Engineering Effective Responses to Influenza Outbreaks

Stan N. Finkelstein
Engineering Systems Division, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; and
Harvard Medical School, Boston, Massachusetts 02115, snf@mit.edu

Richard C. Larson
Engineering Systems Division, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, rclarson@mit.edu

Karima Nigmatulina
Institute of Master Planning for the City of Moscow, Moscow 125047, Russia
knigmatulina@genplanmos.ru

Anna Teytelman
Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, teytanna@gmail.com

We present a policy-oriented summary of our six-year "service-systems-focused" research into pandemic influenza. We cover three topics: (1) R_0, the basic reproductive number for the flu; (2) NPIs, non-pharmaceutical inventions to reduce the chance of becoming infected; and (3) flu vaccine allocations. We use a service-systems framing and mathematical modeling approach incorporating theories and data on the spread and control of influenza. We examine how behavioral actions and governmental policies, thoughtfully derived, can minimize influenza's societal impact. There is widespread misinterpretation that R_0 is a numerical constant of a given virus. We argue that it is not, but rather that its value is largely determined by local conditions and actions, many under our individual and collective control. This control is, in the absence of vaccine, intelligent use of NPIs—highly effective in reducing the spread of influenza. Our vaccine analysis relies on government data depicting flu-like cases and vaccines administered during the 2009 H1N1 outbreak. During that outbreak, barely half of all states received allotments of vaccine in time to protect any citizens. The method of vaccine deployment—in proportion to census population—ignored the temporally uneven flu wave progression across the United States.

Keywords: healthcare; influenza; non-pharmaceutical interventions; health service systems; health policy
History: Received January 20, 2015; Received in final revised form April 1, 2015; Accepted April 17, 2015.

1. Introduction

Seasonal outbreaks of influenza are costly in human and economic terms. In the United States, tens of thousands typically die each year, with orders of magnitude more worldwide. Annual economic losses are tens of billions of dollars (McKibbin 2006). Pandemics, occurring an average of three times per century (Sandman 2007), have the potential to be more disastrous than an exchange between warring nations. The 1918–1919 "Spanish Flu" reportedly killed at least 40 million people worldwide, more than the number of deaths caused by World War I (Flu.gov 2010).

However, today we are not impotent in the presence of the flu. As individuals, families, workplaces, and governments, we have the ability to greatly influence the severity and magnitude of any given flu outbreak. This is true even without available vaccines, as was the case during much of the rapidly spreading 2009 H1N1 pandemic, and also in 2014–2015 when the seasonal vaccine offered only limited protection to the prevalent H3N2 strain.

Here we review from a behavioral and decision perspective the key findings of our six years of research on influenza. We focus on three key interrelated topics: R_0, the *basic reproductive number* associated with influenza spread through a population; non-pharmaceutical interventions (NPIs), relating to changes in personal and societal behavior that can reduce the chance of becoming infected; and vaccine, the medical intervention that significantly reduces the chance of becoming infected. Using simple axiomatic mathematical models in conjunction with newly available data as well as established literature, we believe we offer fresh insights that are important for both individuals and governments charged with effective system response to influenza. We believe that insights gained from historical events that might be viewed as "natural experiments" offer validation of the concept that both vaccines and NPIs, properly implemented, acting separately or together can mitigate the course of an illness outbreak.

An interdisciplinary team, involving operations researchers, physicians, systems engineers, policy analysts, management specialists, and a historian (John M. Barry), has carried out the work. Our focus—employing the

methods of service science and health systems analysis broadly defined (Rouse and Cortise 2010)—is on flu illness *prevention*, not health*care*, but maintaining wellness. The effort involved close collaboration with research colleagues at the Harvard T.H. Chan School of Public Health, and we are grateful for their cooperation and helpful feedback throughout our multiyear effort.

2. The Reproductive Number R_0

The most critical parameter—used in almost everyone's modeling of influenza transmission—is R_0, the *basic reproductive number* or *basic reproductive ratio*. The number R_0 is usually defined to be the average number of new infections generated by a "typical" newly infected person in a population of 100% susceptible individuals. The concept is borrowed from population studies, originating in Germany in the 1880s, where R_0 was defined to be the mean number of baby girls that a newly born baby girl would have in her lifetime. In this context, an R_0 value greater than 1.0 indicates a growing population, whereas a value less than 1.0 indicates a shrinking one. With infectious diseases, R_0 has comparable interpretations. Early in a flu outbreak, R_0 is the generation-to-generation average growth factor of newly infected individuals, "new births" of influenza cases. An R_0 of 2, for instance, indicates that early in the outbreak a typical newly-infected person transmits the disease to an average of two others. Note that this is an average, and any given infected individual may infect some number different from the average. We can see that an R_0 value greater than 1.0 causes a near-term exponential increase in the generation-to-generation number of infected individuals. This is the true interpretation of "going viral". An R_0 value less than 1.0 virtually guarantees that the illness will die out.

In population studies, the value of R_0 depends on many factors—age of marriage, a couple's desire to have children, use of birth control, etc. In 1879 in Germany, R_0 was estimated to be 1.06, indicating a stable and slowly growing population (Heesterbeek 2002). Today it is estimated to be 0.70, indicative of "negative population growth" (Larson and Nigmatulina 2009). Therefore, we see that in demography, the value of R_0 depends on personal behavior. We will argue that the same is true for infectious diseases such as influenza.

How are numerical values for R_0 determined? The World Health Organization (WHO) and U.S. Centers for Disease Control and Prevention (CDC) often "announce" the R_0 value for a new flu virus. For instance, the numerical values estimated for R_0 for the 2009 H1N1 flu tended to be between 1.4 and 1.6, with some estimates above 2.0 (Prachayangprecha et al. 2010). These announcements and associated research papers seem to assume a world in which R_0 is beyond human control, as if R_0 were like the mathematical constant π, a type of "nature's constant" attributed to a given virus. But common sense and simple mathematical reasoning suggests that R_0 is not at all a constant of nature. It is to a large extent under our individual and collective control.

The common sense part is straightforward: Just consider extreme forms of social distancing during a flu season, say where all work is done at home via the Internet and telephone, all children are homeschooled, and groceries and other essentials are delivered to one's door in a 100% antiseptic way. Such extreme behavioral change would surely reduce R_0 compared to our usual face-to-face ways of carrying out our day-to-day lives. We are not suggesting such extreme behavioral change, but offer it as an illustration of how behaviors can affect R_0, thus altering the spread of the disease.

Our discussion of R_0 proceeds from simple to more complex. Our perspective is conceptual, not statistical, although later we discuss historical cases that support our cause-and-effect behavioral assumptions and policy conclusions. Our point of view is that limiting the number of close contacts between infectious and susceptible persons will reduce flu prevalence. It is well established that flu viruses can easily be passed from an infected person to a susceptible person during a handshake, hug, kiss, or similar close encounter. Breathing the air of a nearby sneeze or cough can also pass the infection. Our approach is analogous to that of Kaplan and O'Keefe (1993) in reducing usage of dirty needles by intravenous drug users in order to reduce incidence of HIV infections. It is also similar to assumptions of others that reduction of incidence of unprotected sexual encounters will reduce STDs including HIV.

We start with a homogeneous fully-mixing population, as is assumed in many traditional S-I-R (susceptible, infected, recovered) models and then add complexity reflecting heterogeneous populations. Our approach is also a S-I-R model, only we use probability models rather than differential equations to advance people from state to state—susceptible, infected, recovered. We develop a simple equation for relating extent of face-to-face contact and R_0. The relationship (Larson and Nigmatulina 2009) involves an infectious person's frequency of human-to-human contact (λ) while infectious and his or her conditional probability of transmitting illness (p), given "close contact" with a susceptible individual.

Suppose I come face-to-face with N people on a day that I am infectious but asymptomatic. Most people who become infected with the flu have one such day before they feel and appear sick, and not being able to identify these people is what makes eradication of the flu so difficult. Define an "indicator variable" as follows:

$$X_i = \begin{cases} 1 & \text{if person } i \text{ becomes sick as a result of exposure to me,} \\ 0 & \text{if person } i \text{ does not become sick as a result of exposure to me.} \end{cases}$$

Now, we let NI be defined as the number of people I will infect on this infectious day. NI can be written as simply counting the indicator variables:

$$NI = X_1 + X_2 + X_3 + \cdots = \sum_{i=1}^{N} X_i.$$

Suppose, for example, $N = 50$ and that all X_i's are 0 except for X_9, X_{18}, and X_{45}, each being equal to 1. In that case, I have infected three of the 50 individuals I have come face–to-face with on this day.

Now, at any given level of intensity of face-to-face contact, there is a probability p that I will pass the infection on to the person I am facing. Using this fact, we can write an expression for the mean number of people I will infect on this day. It is simply the mean of $NI = X_1 + X_2 + X_3 + \cdots = \sum_{i=1}^{N} X_i$, which equals Np. We thus have a simple expression for R_0:

$$R_0 = Np.$$

In practice, beforehand we do not know the numerical value for N, only the average number of daily contacts. In our final equation we replace N with λ, the *mean* number of people we interact with on a given day, obtaining

$$R_0 = \lambda p. \tag{1}$$

This simple equation shows that R_0 is a function of both the inherent properties of the given virus—as represented in part by p—and the population's behavioral responses to it—as represented by both λ and p. If the duration of asymptomatic infectiousness is more than one day, then the rate parameter λ in Equation (1) is redefined to be the mean number of human contacts during the asymptomatic infectious period. Designing an empirical study to confirm Equation (1) would be difficult, if not impossible. However, we believe strongly that our conceptual perspective is supported by numerous historical examples that demonstrate that reducing personal contacts via social distancing and hygiene measures is statistically associated with reductions in influenza transmission.

As an illustrative application of Equation (1), consider a checkout clerk at a grocery store who serves an average of 100 customers per day. Suppose that this clerk became infected during the previous weekend, and he starts his weekly daily work pattern again on Monday. A significant issue with the flu is that an infected person can be infectious and asymptomatic for a day or more, unknowingly spreading the flu to others, and only later come down with flu symptoms—at which point he or she will most likely self-isolate and usually eventually recover. This fact makes R_0 larger than it would have been, had there been no asymptomatic infectious period prior to arrival of flu symptoms. Our store clerk is infected already when he works on Monday, but not yet infectious. No customer he serves on Monday becomes infected from him. He becomes infectious Monday evening and works all day Tuesday, infectious but asymptomatic. Given the clerk's lack of frequent hand washing, together with his touching his face and mouth (now with infectious virus), assume that each customer he serves on Tuesday has a 10% chance of becoming infected by the clerk. Then, applying Equation (1), the clerk's value for R_0, the mean number of individuals he will infect, is $100 \cdot (0.10) = 10$. Of course, if Tuesday happened to be a slow day with only 50 customers, he would have infected only half as many people: $50 \cdot (0.10) = 5$. These are mean values, and on any given day the actual numbers infected may be different from the mean. We can close this example by assuming our clerk begins to develop flu-like symptoms Tuesday evening, and then decides to stay home and totally recover before returning to work.

Since R_0 is proportional to the product of p and λ, we need to ask: *how do we reduce λ and p*? One reduces λ simply by having fewer face-to-face contacts each day. Our extreme social distancing example illustrated this, as did our grocery clerk example. What about p? Vaccines are probably the best-known way of reducing p. But p is a function of three factors: (1) innate infectivity of the virus; (2) hygienic practices of the infected and infectious individual having the close contact; and (3) the vaccine status and hygienic practices of the noninfected but possibly susceptible individual on the other side of the close contact. Whereas (1) is uncontrollable, (2) and (3) represent opportunities to reduce the likelihood of passing the infection along to the individual who is susceptible without vaccine. Helpful personal hygiene practices include frequent aggressive hot-water hand washing and

not touching one's face with one's hands. Localized social distancing also plays a critical role, as in avoiding handshakes, perhaps bowing or elbow bumping instead. In our models, R_0 is the key measure incorporating the range of determinants of the spread and control of an infectious outbreak.

As an illustrative example, suppose we find ourselves living in a flu wave that is said by authorities to have a high R_0 value of 2.0. That means, from Equation (1), $R_0 = \lambda p = 2.0$. If we can collectively reduce our λ by 50% or p by 50%, we will have reduced R_0 to 1.0, the value indicating no further growth. Any additional reduction will mean the flu wave will die out, perhaps rapidly. But our task is even easier, as any *combined* reduction of both λ and p simultaneously that reduces their product by 50% will yield the same good result. For example, reduce λ by 30% and p by 30%. Then $(0.7)\lambda \cdot (0.7)p = 0.49\lambda p = (0.49) \cdot 2.0 = 0.98$, or slightly less than 1.0, the desired result. The key here is that, individually and collectively, we have some control over R_0. And most flus have reported R_0 values of less than 2.0, typically 1.2–1.6, and thus our work required to reduce λ and p, and thus R_0, is even less.

The number R_0 is the mean of new infections generated by a newly infected individual very early in the outbreak, when nearly everyone is susceptible to the illness. But as the disease progresses, those infected early presumably recover and re-enter the active population but are now immune to additional infection. As those who are immune resume their usual daily activities, a fraction of the "λ-frequency" human contacts with infectious individuals will have no effect on transmitting the infection, as the interaction will be with an immune person. Therefore, the mean number of new infections generated by a newly-infected person will decline as the flu progresses, as an ever-growing fraction of daily human interactions is with now-immune individuals. We need to define $R(t)$ as the mean number of new infections caused by a "typical" newly-infected person during generation t of disease progression, $t = 1, 2, 3, \ldots$. A generation of the flu is typically 2.5 days. Because more and more people recover and become immune to infection as the disease progresses, we always have $R(t)$ becoming smaller with each generation, i.e., $R_0 \geq R(t) \geq R(t+1)$, $t = 1, 2, 3, \ldots$,. That is, the rate of growth slows and eventually stops growing when $R(t) = 1$, and then declines. As an example, R_0 could start out as 2.0. For the next several generations, where virtually everyone in the population is still susceptible, $R(t)$ will remain near the R_0 value of 2.0, decreasing only imperceptibly. This creates early exponential increase in the number of new cases, generation by generation. Eventually, a measurable fraction of the population is recovered and immune to further infection, and then $R(t)$ begins to show measurable decline.

In a sense, $R(t)$ equaling 1.0 is a tipping point, the moment when there is no further growth in the generation-to-generation number of people infected. We need to ask: what fraction of the population must be immune at the tipping point? Suppose at the tipping point we let f_S be the fraction of the population that is still susceptible, and $f_I = (1 - f_S)$ be the fraction that is immune. Then, at this point, we can rewrite Equation (1) as follows:

$$R(t) = 1.0 = \lambda p f_S = R_0 f_S, \tag{2}$$

or

$$f_S = 1/R_0, \tag{3}$$

implying that the fraction f_I that must be immune is

$$f_I = 1 - f_S = 1 - 1/R_0. \tag{4}$$

Here, at the point when the fraction of the population that is immune reaches $f_I = 1 - 1/R_0$, we achieve "herd immunity." Personal immunity can be achieved in at least two ways: recovery from the flu infection (meaning one's body now has immunity-creating antibodies) or vaccine. When a population has achieved herd immunity, there can be no further generation-to-generation growth in the disease. As an example, if $R_0 = 2.0$, then herd immunity is achieved when 50% of the population becomes immune. If, as is more typical with the flu, $R_0 = 1.4$, then herd immunity is achieved when $f_I = 1 - 1/R_0 = 1 - 1/1.4 = 28.6\%$ of the population becomes immune. If one adds the option of vaccines, thereby adding to the fraction immune, we see that 28.6% is an achievable goal.

We need to point out that all of our discussion of R_0, $R(t)$, and herd immunity to this point has assumed a homogeneous fully-mixing population with each person behaving similarly. Unfortunately, life is not so simple.

Stochasticity. Although R_0 provides an easy and computationally intuitive basis for describing disease dynamics, it has a number of limitations that tend to be distributional. The value of R_0 is an *average*, implying that it derives from a probability mass function (pmf) whose mean is R_0. The random variable associated with the pmf is the number of individuals that a random newly-infected person infects. Consider an R_0 value of 2.0. At one extreme, all the probability mass may be located at 2, and with deterministic regularity each newly-infected

person early in the pandemic would infect exactly two others. But that is utterly simplistic, as no human population would behave in such a robotic manner, no population giving birth to babies, and no population spreading the flu. Perhaps the pmf might have a geometric shape, starting at 0 and having mean 2. With such a geometric pmf, 33% of newly-infected individuals would infect no one else, whereas about the same percentage would infect three or more. This model depicts a great deal of variability on the number of new infections generated by any newly-infected individual. One can think of limiting cases at an extreme with most probability located at zero and with a small mass at a large number such as 40, still with mean equal to 2. This type of situation involves so-called "super-spreaders" who if active early in the pandemic can catapult it to major status but if they do not appear early will result in a flu that dies out rapidly, even with an R_0 value of 2. We must remember that R_0 is the mean of a probability distribution, and its variance and, in fact, entire distribution will play a role in the evolution of the disease.

Heterogeneity. In addition to stochasticity, we have another distributional issue with R_0—heterogeneity of the population. Members of a population are heterogeneous with respect to their personal characteristics and behaviors. In a sense, each person in the population has his or her "own" value of R_0, like our store clerk, which is to a large extent under the control of the individual.

As an outbreak of influenza evolves, public health decision makers receive aggregate statistics in the form of the number of people reporting to physicians with flu-like symptoms, related hospital admissions, flu-related deaths, and vaccinations administered. Yet aggregate statistics hide the fact that early transmission and propagation of the disease are driven largely by particular segments of the population: (1) those who are highly active in daily face-to-face encounters; (2) those who are overly prone to become infected given exposure; and (3) those who shed virus and spread the disease more than average. Any person can be characterized along a spectrum of these three attributes: social activity, proneness to infection, and proneness to shed virus and spread infection. Those who are at the "right-hand-tails" of one or more of these distributional attributes play a significant role in the early spread of the disease. Such individuals, due to early infection and later immunity, drop out of the susceptible population near the middle and almost certainly by the end of the outbreak. An unfortunate example of a critical subpopulation that becomes infected early is healthcare workers, whose daily job puts them in close contact with infected and infectious individuals.

The best available data on variability of human contacts is from two published sources. Fu (2005, 2007) reports a study of 3,000 respondents from nine countries and 46 different settings who were asked to estimate daily personal contacts, including face-to-face, telephone, mail, and Internet. For Taiwan, Fu found that 83% of reported contacts were face-to-face (Fu 2005, 2007). Fu's data show a significant right-hand tail of the distribution, with 28% of the respondents reporting fewer than 10 personal contacts per day and 11% reporting more than 100. In a separate report, the Mossong group conducted a thorough study of contacts by participants in eight European countries (Mossong et al. 2008). Participants were asked to record their daily contacts, defined as either "skin-to-skin" or a two-way conversation with three or more words in the physical presence of another person. The information from participants' diaries was weighed to match the demographics of participating countries. The group published distributions of daily contacts by individuals from each country. We incorporated data from four of those countries, Belgium, Great Britain, Germany, and Poland, into our models and subsequent analyses. In comparing countries, the R_0 values implied by the data varied by almost a factor of two, with Germany's being the lowest R_0 at 1.33 and Poland the highest at 2.54.

Population heterogeneity—due to widely differing rates of social contacts and also due to infection proneness and virus shedding behaviors—plays a key role in the speed of infection spread. Everything else being equal, it is reasonable to assume that people who are most socially active are more likely to spread it to others. Those who are socially active *and* susceptible to infection are even more likely to spread it to others. Those who are socially active *and* susceptible to infection *and* "efficient virus shedders" are the most efficient virus spreaders. To understand the dynamics of flu spread, or the spread of any human-to-human infectious disease, one must account for such population heterogeneities.

The three dimensions of heterogeneity bring into question the estimation and even the very definition of R_0. People who have one or more of these attributes—socially active, infection-prone, and/or efficient virus shedders—largely drive the early exponential growth of the disease. This suggests that our definition of R_0, if it is to represent the generation-to-generation early exponential growth of the disease, needs to be more nuanced than simply: ". . . average number of new infections generated by a '*typical*' newly infected person in a population of 100% susceptible individuals." "Typical" is too vague. One suggested change is to replace "typical infected person" with "typical face-to-face interaction with an infected person." Such a change would automatically account for those early in the disease growth with greater-than-average "scores" in our three dimensions–focusing on interactions and not individuals. As an example, from Fu's (2005, 2007) data, a random person from the 11%

of the population having more than 100 human interactions per day is much more likely to spread infection than a random person in the 28% having fewer than 10 human contacts per day. We could have a situation in which the majority of the members of the population have "personal R_0 values" less than 1.0 while a minority fraction drives the societal R_0 to a value significantly greater than 1.0. In other words, we could average the personal R_0 values of a nation's "typical people" and find the population-averaged R_0 to be less than 1.0, yet the nation-averaged *experienced* R_0 significantly greater than 1.0—for the reasons we have just outlined.

The most common models for influenza spread follow some variant of the S-I-R compartmental approach, where each person is *s*usceptible, *i*nfected, or *r*ecovered (or deceased). These models are most often used in a homogeneous setting, where all people in a compartment behave identically and mix randomly. One may call the approach "models of statistical clones." Teytelman and Larson (2012) generalize that approach to eliminate the need for a finite number of groups—classes of statistically identical individuals, and instead, introduce a continuous distribution for all the key parameters in question, in essence employing an infinite number of classes. Their generalized model deals with all three attributes previously introduced: social activity, proneness to infection, and proneness to spread infection. The model relies on just a few equations that define the state of infection at a given time.

Our approach has been to use discrete-time models and account for heterogeneity via proportional mixing, where an individual is likely to become infected in proportion to his or her contact rate. We introduce a continuous distribution for three parameters of interest—social activity, proneness to infection, and proneness to spread infection. The initial focus is on contact rates, the available measure of social activity. The unit of time is a generation of influenza, defined here as the two-to-three day period during which a person becomes infected and soon infectious and interacts in society. The model relies on difference equations that define the state of infection at a given time and allow the calculation of $R(t)$, the analog of R_0 at any point in time, as the outbreak evolves, as discussed earlier. The full formulation of the model has been published elsewhere (Teytelman and Larson 2012). Future research should focus on obtaining reliable data providing information on the joint distribution of our three attributes.

3. NPIs: Human Contact and Behavior

When studying and modeling sexually transmitted diseases, especially HIV/AIDS, behavioral changes are often cited as the main factors determining transmission dynamics, but when it comes to modeling flu, behavior is almost always ignored. This is puzzling, and—in our opinion—quite incorrect. Few would dispute the observation that people alter their behavior during an outbreak by adopting more diligent hygiene, and by decreasing their frequency and intensity or closeness of human contacts. Recent history has provided us with multiple examples of people responding to news of a disease by altering their daily behavior.

Consider behavioral changes that occurred during SARS in 2003 in Hong Kong. One survey indicates that during the SARS outbreak, 87% of the Hong Kong residents covered their mouths while sneezing or coughing, 76% of individuals wore masks, 65% washed their hands after contact with possibly contaminated objects, and more than 50% used diluted bleach for household cleaning. There was a sharp reduction in many discretionary activities such as attending social events, shopping, and going to restaurants. Residents who thought that they might have been exposed to SARS isolated themselves voluntarily for up to 10 days. The Hong Kong Government ordered closings of schools, libraries, swimming pools, and Kowloon Bay Sports Center. Hong Kong tourism was crippled in March 2003 when the World Health Organization (WHO) issued a rare warning for travelers to avoid Hong Kong and China's Guangdong Province. As a result of weakening demand, airlines slashed more than a third of flights, and hotels in Hong Kong reportedly were up to 90% empty.

SARS was stopped, and yet no pharmaceutical cure was found (Lo et al. 2005). There is additional evidence of the beneficial effects of hygienic steps and social distancing. The incidence of other acute respiratory viral diseases (seasonal influenza, parainfluenza, respiratory syncytial virus, and adenovirus) during the key months April and May 2003 dropped 90% compared to seasonal norms. This is additional best evidence that behavioral modifications can dramatically reduce the spread of respiratory infections.

To the best of our knowledge, the eradication of SARS was due to collective behavioral changes of the overall population and medical caregivers, in effect causing R_0 to drop significantly below 1.0. This represents an existence proof that R_0 can be largely determined by individual and collective behavioral change. This is a profound result. It suggests that R_0 is not defined in the abstract as a constant of any given infectious disease.

John Barry, in his classic book *The Great Influenza* (Barry 2004), reports on the results of bad social distancing and behavioral changes during the 1918–1919 pandemic influenza. He cites Philadelphia as a worst-case scenario,

due to authorities there holding a "Liberty Loan parade" at the height of the local flu wave, an event that acted as an efficient concentrator and accelerator for the flu (Barry 2004, pp. 208–209). Others (Bootsma and Ferguson 2007, Hatchett et al. 2007) cite St. Louis as one city that significantly knocked down its flu wave by imposing restrictions not dissimilar to those of Hong Kong many decades later. The trouble with St. Louis is that authorities there celebratorily removed the social distancing restrictions too early, and the St. Louis flu wave resumed its upward trajectory. This resulted in a rare "double-humped" flu wave for the city. More recently, Chowell et al. (2011) studied progression of the H1N1 pandemic in Mexico in 2009. Their findings support the theory that social distancing reduces incidence of infection:

> "We estimate that the 18-day period of mandatory school closures and other social distancing measures implemented in the greater Mexico City area was associated with a 29%–37% reduction in influenza transmission in spring 2009. In addition, an increase in R was observed in late May and early June in the southeast states, after mandatory school suspension resumed and before summer vacation started."

All of these cases provide strong evidence supporting the relationship between human behavior and transmission of the disease.

Given the connection between human behavior and disease transmission, it makes no sense for the WHO or CDC to state publicly that a new influenza virus is circulating the Earth with an R_0 value of, say, 1.432. Rather, the local population and their individual and collective behaviors contextually determine R_0. In the future, it is entirely plausible that when a novel virus surfaces there will be communities for which R_0 is less than 1.0 and other communities, such as those living in close and closed quarters, where R_0 could exceed 2 or 3 or more. Any modeling analysis that ignores behavioral changes and local context removes our greatest disease-progression control strategies.

Recognizing the importance of human behavior and the likelihood of becoming infected due to "high λ" and/or "high p" (Equation (1)), we looked at recurring instances where people could be exposed to such greater-than-average risks. Healthcare workers (HCWs) are a prime example. Their reported use of flu vaccine in recent years has ranged from 30% to 62%, this maximum during the 2009–2010 H1N1 pandemic (Stewart et al. 2011, see also National Foundation for Infectious Diseases 2008). From the perspective of Equation (1), relating flu infections to frequency of daily contacts and conditional probability of infection given contact with an infectious individual, we cannot think of a more important class of people to vaccinate than HCWs. In considering Equation (1), a susceptible uninfected healthcare worker serves say N flu-infectious patients per day. Suppose that the probability of becoming infected from any one of the patients is p, and that the chances of infection are independent from patient to patient. Then, after treating N patients, the probability that the healthcare worker is now infected is 1 minus the probability that no patient infected him or her, namely

$$P_I(N) = 1 - (1-p)^N. \tag{5}$$

For any positive value of p, the issue of infection is not *if*, but *when*. If $p = 0.10$ (or 10% for each patient), then after serving 50 patients each infectious with the flu, the probability that our healthcare worker is now infected is

$$P_I(N) = 1 - (1-p)^N = 1 - (0.9)^{50} = 1 - 0.00515 = 0.995,$$

in other words, infection is a virtual certainty. If our HCW were to wear gloves, a mask, and take care not to touch surfaces and his or her own face, the probability p may be brought down to, say, 0.02 (namely a 2% chance of infection per patient). In that case, after serving 50 patients, our healthcare worker has a probability of being infected from one or more of the 50 patients equal to

$$P_I(N) = 1 - (1-p)^N = 1 - (0.98)^{50} = 1 - 0.364 = 0.636,$$

which is still incredibly high. According to (Stewart et al. 2011), "During an average season, 23% of HCWs are infected with the virus, show mild symptoms, and continue to work despite being infectious." While the reported 23% is less than our model with 50 patients, the possibility for passing infection along to others remains too high. Thus we see the need to utilize vaccines for each and every healthcare worker. The scenario is much worse than we have depicted, because once the healthcare worker is infected, he or she soon becomes infectious asymptomatically for at least one day, a day in which many of his or her patients are hospitalized patients with non-flu serious health conditions with compromised immune systems. Should one of these patients become infected with the flu, they may suffer severe complications, perhaps leading to death. The lesson here is clear: HCWs need to be vaccinated against the flu.

Table 1. Suggested Home Care Flu Kit

Things to get:	Typical cost
Detergent-based soap	$10
Alcohol-based hand sanitizer	$10
Window fan	$40
High efficiency particulate air (HEPA) filter system	$0–$600
Ultraviolet light unit (some with HEPA filters)	$180–$370
Tissues	$5
Face masks (25)	$10
Approximate total	$250–$1,000

Those at home caring for a loved one sick in bed with the flu are also healthcare workers, and unless they are careful, they are prime candidates for infection—due to virus-containing air droplets from coughs and sneezes of the at-home patient or from touching surfaces that were touched by the patient (Atkinson and Wein 2008, Wein and Atkinson 2009, Cauchemez et al. 2009). We dedicated a fraction of our flu research not on mathematical models but in the literature pertaining to ameliorative steps one can take—at home and in hospitals—to reduce the chance of such close contact infections. We reviewed the post-2000 scientific literature to identify steps that in-home caregivers can take to reduce the chances that they and other household members will become infected in the home. The literature was fragmented and disparate. We attempted to bring the key results together in one coherent package. The results of our scientific literature review can be found in Finkelstein et al. (2011b). The suggested components of an at-home flu "kit" are shown in Table 1. For those with more flexible budgets, we argue the benefits of adding a space heater (about $50) and a humidifier (about $100).

In conjunction with purchasing and installing the kit, here are the things to do and not do:

- Wash hands frequently (particularly after shaking hands with people or coughing) with hot soap and water for at least 30 seconds, and dry.
- Try not to touch your face with your hands.
- Practice careful bathroom etiquette.
- Avoid direct hand contact with surfaces that are likely to be contaminated.
- Hold meetings via telephone and email when possible.
- Read and study http://www.ifh-homehygiene.org/2003/index.html.

The paper by Finkelstein et al. (2011a) was designated for AMA PRA Category 1 Credit™. Physicians and other health professionals earn credit by reading the article in the online issue of DMPHP and taking a quiz online. The behavioral recommendations from the paper have been empirically verified in a controlled experiment in China (Zhang et al. 2013). As of this writing, the CDC is undertaking a trial implementation of the key recommendations of the paper with CDC employees—to reduce their chance of getting the flu when they are caring for a flu-infected loved one at home.

4. Disease Dynamics and Vaccine Distribution

To this point we have focused on R_0 and how its value can be reduced by individual and collective behavioral changes. These are our control strategies in the absence of pharmaceutical intervention. Now we switch to flu vaccines, its distribution, and administration. Being vaccinated against the seasonal flu is our second major behavioral choice in our arsenal of flu-preventing tools.

Getting a flu shot is not only good for the individual but also for the community. If an individual becomes infected with the flu, he or she may pass it on to others, up to an average of R_0 additional people if the infection occurs early in the flu wave. The decision to accept the flu vaccine is a type of "reverse tragedy of the commons" (Hardin 1968). In the original tragedy of the commons, one pictures a number of farmers who each place an increasing number of cows onto a "commons" grazing field. At the beginning, each cow remains healthy and produces lots of milk. But eventually, as each individual farmer is trying to maximize his own narrow self-interests by adding additional cows, the commons grass becomes over-grazed, the cows start to lose weight and give less milk. The process can deteriorate to total tragedy, ultimately with the cows dying. In a reverse tragedy of the commons, one individual's act can help not only that person but also the broader community. Our modeling analyses suggest that an individual getting a flu shot before the flu wave hits a fully susceptible population could typically prevent an average of 1.9 flu infections (Larson and Teytelman 2012). Therefore, the vaccine in one individual reduces, in an average sense, flu infections of others as well. Each individual's flu shot gets us all a bit closer to the desired state: herd immunity.

Vaccines for seasonal flus are almost always available in abundance before the flu season starts. These vaccines have their own problems, as those creating the vaccine typically one year in advance must "guess" the likely dominant flu virus strain during the upcoming flu season. They create a vaccine that covers typically three different but plausible strains, often using sophisticated statistical techniques (Wu et al. 2005). Then ultimately the flu season starts and we all see how closely the season's flu vaccine matches this year's dominant strain. For instance as we write this in the winter of 2014–2015, the dominant strain this year in the United States is a variant of Influenza A (H3N2) and is not well covered by this year's vaccine. Evidence suggests it is marginally effective, only 23% effective according to the CDC (Edney 2015). Therefore, a flu vaccine, even if available in abundance, does not shrink to zero the probability of an individual becoming infected if exposed to the virus. Although a vaccinated person who comes down with the flu usually experiences a milder case than someone who is not vaccinated, the fact that the vaccine does not provide 100% protection should motivate us all to follow the NPI steps outlined earlier.

The "pandemic" outbreak of 2009 H1N1 influenza was different. During that flu season, any anticipated seasonal flus were replaced with a novel H1N1 virus, where "novel" implies that nearly everyone is susceptible to the illness and that no vaccine exists. The H1N1 flu waves started in earnest in the southeastern states of the U.S. during August 2009, and vaccines did not become available until late October, and then only in limited quantities. We applied our modeling approach to this problem to better understand the health effects of delayed vaccine production and distribution and perhaps also to identify improved ways of distributing vaccines when arriving late in the flu season.

Using 2009 data from the CDC and state health departments, we estimated the influenza-like-illness (ILI) epidemic curves for the United States as a whole and for 48 states. The CDC considers ILI data to be effective means of following the dynamics of progression of the outbreak. Sentinel sites report the proportion of outpatient visits, hospitalizations, and deaths associated with ILIs to the CDC via ILINet, an online reporting system (CDC 2010a). The CDC tabulates these data on national and regional levels and publishes results weekly in *FluView* (CDC 2010b). We compared the epidemic curves we derived with two sources of vaccine distribution data. The first is vaccine shipment data, which track, for all 50 states, the number of doses of vaccine shipped to each state over time. The second source provided data on vaccines actually administered, as each healthcare provider was required to report numbers of flu vaccinations administered by state and local health authorities before being given additional vaccines. We obtained this latter information from individual health departments of nine states (Finkelstein et al. 2011b).

4.1. Results

We found that in 24 of 50 states the outbreak had already begun to decline before any individuals were protected by vaccination immunization. Further, among 11 states, no more than 2% of the state's residents were vaccinated before the outbreak had peaked.

For each of the 11 states, our model was fitted to the reported ILI data to create two separate model-estimated epidemic curves: the first assuming no vaccines delivered and the second incorporating actual vaccine administration data for the state. We also generated a third model-based epidemic curve showing the curve if the vaccine had been delivered two weeks earlier than actual. We were then able to infer the number of infections that were averted due to the administration of vaccine, even if late; and the number of infections that would have been averted if the vaccine supplies had been received two weeks earlier. Averted infections ranged from as much as nearly 14% of the population in Massachusetts, where the outbreak occurred later, to as little as 0.14% in Mississippi, which experienced a much-earlier outbreak.

4.2. Vaccine Allocation

In the 2009 H1N1 pandemic, as stated earlier, vaccines arrived late and in limited supply. This is not surprising due to the six-to-twelve month delay between identifying a novel flu virus, inventing an appropriate new vaccine, and manufacturing it for distribution. The pandemic was already well underway in the United States when vaccine distribution commenced in October 2009. Early deliveries were rationed and delivered to states by the CDC in direct proportion to each state's census population, regardless of the status of the flu wave in the state. This deployment method is at least partly driven by perceptions of equity and other "political" considerations. For example, if Alabama has 1.7% of the U.S. population, Alabama citizens and taxpayers may think that 1.7% of manufactured vaccines is "theirs."

The fundamental problem with the CDC vaccine allocation method is that many shipments arrive long after they would have been beneficial. In 2009, for states in the Southeast, the October-arriving vaccine shipments were analogous to "closing the barn doors after all the horses had left." That is, there were virtually no more

people to vaccinate who were otherwise still likely to become infected with the flu. The flu waves in the Southeast had run their course, starting in August when public schools opened. Yet, while the Southeast was still getting proportional shipments of vaccine, northern and western states were in dire need of additional doses, as the flu waves there had not yet peaked, and many susceptible individuals could have been immunized.

Our model results suggest that the CDC's population-based flu allocation approach is far from optimal, as it does not focus on what we believe to be the critical performance measure: minimize the total number of flu infections that will occur nationally. Rather its objective is to equalize per capita distribution of the vaccine regardless of its potential flu-averting benefits nationally.

A population-based vaccine allocation policy has a fire brigade analogy. Imagine a fire brigade delivering buckets of water to fire fighters attempting to stem fires in a row of houses. At a given time, some of the houses have been totally destroyed by fire, some untouched, and others are in various states of growing active fire damage. The CDC method, applied to fire-fighting water allocation, might allocate water to houses in direct proportion to the square footage of each respective house, regardless of its current damage status: totally burned down, untouched, or in some intermediate state of active fire damage. While this clearly makes no sense in fighting fires, vaccines delivered late during a flu season are directly analogous to water fighting fires, instead here we have vaccines fighting flu infections, and various states in the U.S. are "totally burned down" (i.e., the flu waves have come and gone), untouched (i.e., the flu waves have not yet started), and in an intermediate state with flu waves rising.

A better policy would be to deploy vaccines not in proportion to state populations, but to vulnerable regions that have seen fewer cases, that will have a higher fraction of its population susceptible, and thus where a vaccine can avert the maximum number of future infections. That is, throw the water now on those houses that will benefit most from the treatment. Using our mathematical models, we have developed an adaptive vaccine deployment method that focuses weekly vaccine distribution on those now-active states where immunization can avert as many infections as possible. States untouched by flu can wait; states having completed flu waves will not benefit from additional vaccine. Our flu vaccine deployment method, if it had been used in 2009, would—we estimate—have averted about 5,000,000 of the estimated 21,000,000 Americans infected with H1N1 flu (Larson and Teytelman 2012).

As for Alabama wanting "its" 1.7% of vaccine shipments, this year's flu may suggest Alabama gets only 0.5% of the shipments, but next year the state may require 5%, all with the intent of minimizing the total number of infections over the entire country. In fact, as we write this (winter 2014–2015) the regional dynamics of flu progression over the United States are quite different from those of 2009. This year, the seasonal flu appears to have started mainly in the Midwest, with flu wave progression moving both east and west from the center. If Alabama takes a multiyear perspective, sacrificing vaccine doses one year in the cause of a national goal—minimizing the nation's total number of infections—it may mean that over a 10- or 50-year period, Alabama itself would have a minimum possible number of flu infections.

In 2009 the citizens of the Southeast appeared to act in concert with our analysis, in the sense that few individuals accepted the vaccine. In Mississippi, for example, less than 40% of its delivered vaccine was used, most likely due to "flu fatigue." South Carolina managed to immunize only 8% of its population. Had vaccine been available and delivered there before the outbreak peaked (late August and early September), its effectiveness would have been greater with respect to both disease dynamics and participation rate. Similar observations can be made about other states in the Southeastern U.S., where schools opened in August when the flu waves started. With special focus on the difficulties in 2009, see (Rambhia et al. 2010) for a full discussion of getting mass vaccinations in anticipation of pandemic flu.

There are at least two levels of vaccine distribution: federal and state. We have discussed the federal (CDC) method and recommended a fundamental change for late-arriving vaccines during a season of novel flu virus. Such seasons are certain to occur again, we just do not know when. Perhaps equally important are the procedures for allocating federally delivered vaccine doses within a state. Our experience working with states is that current processes are often ad hoc. In considering intrastate vaccine deployments, here too, we must consider the status of the local flu waves within the state, to assure highest marginal benefit from each dose allocated, and also the issues of heterogeneity discussed earlier (i.e., targeting high activity individuals and those who are prone to infection and efficient virus shedders). We believe that fundamental new research is required to invent better two-level vaccine distribution systems. Such systems would aim toward averting the maximum possible number of new flu infections while balancing the need to deploy vaccines to those who, if infected (e.g., the elderly and pregnant women), might suffer the major complications.

Finally, the latest word on flu vaccines: there are scientists who now believe that a universal flu vaccine will be available within five years (McNamee 2015). Should such a vaccine be successfully developed and tested,

one would only need a single vaccination to enjoy lifetime immunity. Under this scenario, all of the logistical problems discussed would disappear.

4.3. Discussion

To engineer an effective health service system response to an outbreak of influenza, one would deploy technology (e.g., vaccine) *and* affect changes in human behavior, to reduce the contact rate and probability of illness transmission. Both NPIs and vaccines are, of course, key components of the public health response.

It is unlikely that society will implement severe measures as they did in 1918–1919 making it "unlawful to cough and sneeze," punishing violators with up to one year in jail. However, even without forceful implementation, people are likely to try to decrease their likelihood of becoming ill by improving hygiene-related behaviors. We control the contact rate, for example, by switching from daily to weekly grocery shopping, or, better yet, having groceries delivered to one's door. If you manage a team of employees, rather than having face-to-face meetings during a flu emergency, have conference calls instead, with many workers telecommuting. Many companies have already created comprehensive pandemic flu plans that include telecommuting, reduced face-to-face encounters, and even increased desk spacing between workers.

Vaccines and NPIs both contribute to reducing the probability that any given face-to-face contact will result in a new infection. Wash hands with hot water and soap several times daily. Do not shake hands during greetings with colleagues. Cough or sneeze into your elbow, not into the open air or your bare hand. Be careful not to touch surfaces that might have recently been contaminated with flu virus. Encourage your city's large employers to stagger work hours, so that public transportation subways and buses are less crowded during now-stretched-out rush hours. Even run the subways and buses with windows opened (Finkelstein et al. 2011a).

4.4. Targeting High Activity Populations

As we have discussed, the at-risk population is heterogeneous in its social activity and susceptibility to contract and transmit illness. Of particular interest is to consider how social behavior influences the propagation of disease. Our many modeling results demonstrate convincingly that targeting high-activity population groups has the greatest impact on how quickly the outbreak can be controlled. High activity members of a population can contribute to mitigating the effects of an outbreak by accepting vaccines to reduce their own susceptibility and transmissibility, by reducing human contacts, and by adopting NPIs to reduce transmissibility.

Vaccines offer the greatest societal benefit when administered early to highly active population members. This observation should be considered when constituting "high risk" groups to be offered early access to immunization. In addition to first responders, healthcare workers, elderly, and chronically ill, a portion of the first available doses of vaccines should be targeted to those individuals having large numbers of daily human contacts.

Diligent personal hygiene among high activity persons benefits not only themselves, but also others with whom they have contact, and can have a disproportionate role in reducing spread. Recall the reverse tragedy of the commons. Hence, there should be great value in targeting these same groups with messaging to adopt NPIs.

How do we target persons with high frequencies of human contacts? Public health practices commonly address school-age children and others spending time in closed and confined quarters. Those who make use of our various transportation networks are also thought of as potentially disproportionate illness transmitters. Consideration has been given to imposing travel restrictions, however, many published articles suggest that it offers low payoff, at best (Nigmatulina and Larson 2009).

Users of all forms of public transportation—subways, trains, buses, and planes—can be considered "high activity" and targeted for behavior change, which can include, depending on outbreak severity, encouragement to get vaccines and practice various NPIs. Short of any mandated shutdown of transport networks, voluntary measures, if adopted, could prove to be major contributions to controlling an outbreak. Potential benefits from this approach extend to individuals engaging in private modes of transportation, including taxicabs, carpools, and even solo commuters, who come in contact with others at retail establishments and fast food restaurants.

The bottom line is that targeting members of a highly active population group to change behavior is likely to be more successful than a broad public campaign. This has proven true in the marketing of many consumer products and in screening for treatable illness, and would be very worthwhile in the control of outbreaks of infectious disease.

4.5. Limitations

Statistician George E. P. Box famously said, "All models are wrong, but some are useful." Our models are no different. And, we tend to use simple models, following Albert Einstein's advice, "Everything should be made as *simple* as possible, *but not* simpler." Mathematical details are found in the cited references authored and co-authored by authors of this paper (see *References* section).

5. Conclusion and Summary Points

Outbreaks of influenza can pose grave threats to lives and the security of our nation—to all nations. Engineering effective response to outbreaks of influenza—seasonal or epidemic—forces us to consider all aspects of the problem logically and systematically. The value of immunization is greatly enhanced when it is deployed in relation to the dynamics of the progression of the illness. The benefits from diligent personal hygiene and social distancing, while widely recognized, can be even greater if public education initiatives were targeted toward population members having high frequency of human contact.

Since we have covered a wide range of topics, we list our key points as follows:

1. A given flu's basic reproductive number R_0 is not a fixed parameter of the flu virus. Its numerical value is to a large extent under our individual and collective control via behavioral changes—social distancing and hygienic steps. We should expect widely varying values for R_0 in different locations and contexts.

2. Even in the absence of an effective vaccine, intelligent implementation of NPIs (non-pharmaceutical interventions) can dramatically reduce incidence of the flu. These NPIs should be a synergistic combination of local (individual, family, and work place) and regional (put in place by government authorities) practices.

3. Any population of humans is heterogeneous with respect to flu infectivity and transmission. People who are prime drivers of the initial flu wave are those who are highly active with many face-to-face contacts each day, and those who are prone to infection (perhaps due to age or compromised immune systems), and those who are efficient virus shedders. They tend to drive the initial exponential increase in the generation-to-generation numbers infected. To the extent that these individuals can be identified, early flu treatments including administering vaccines, should target them (as well as those who—if they were to become infected—could become very ill).

4. When vaccine doses arrive late during the flu season, new methods of vaccine deployment are needed, ones that allocate vaccines to those regions that would enjoy that maximum possible benefit in terms of flu infections averted. Such a system must recognize the two-levels of vaccine allocation: federal and state.

5. New research is needed to develop new models and algorithms, as well as to better understand the role of R_0 in the presence of so much infection process uncertainty and heterogeneity.

We are hopeful that the service science approaches we have described and the results obtained offer the prospect of mitigating future impacts of influenza—seasonal, epidemic, and pandemic.

Acknowledgments

Work on this paper was supported by the Sloan Foundation of New York under a grant entitled, "Decision-Oriented Analysis of Pandemic Flu Preparedness and Response," and under a cooperative agreement with the U.S. Centers for Disease Control and Prevention (CDC) [Grant no. 1 PO1 TP000307-01] "LAMPS (Linking Assessment and Measurement to Performance in PHEP Systems), awarded to the Harvard T.H. Chan School of Public Health Center for Public Health Preparedness (HSPHCPHP) and the Massachusetts Institute of Technology (MIT), Center for Engineering Systems Fundamentals (CESF). The discussion and conclusions in this paper are those of the authors and do not necessarily represent the views of the Sloan Foundation, the CDC, the U.S. Department of Health and Human Services, Harvard, or MIT.

References

Atkinson MP, Wein LM (2008) Quantifying the routes of transmission for pandemic influenza. *Bull. Math. Biol.* 70(3):820–867.

Barry JM (2004) *The Great Influenza* (Penguin Books, New York).

Bootsma MCJ, Ferguson NM (2007) The effect of public health measures on the 1918 influenza pandemic in U.S. cities. *PNAS* 104(18): 7588–7593.

Cauchemez S, Donnelly CA, Reed C, Ghani AC, Fraser C, Kent CK, Finelli L, Ferguson NM (2009) Household transmission of 2009 pandemic influenza A (H1N1) virus in the United States. *New England J. Medicine* 361(27):2619–2627.

Centers for Disease Control and Prevention (2010a) Sentinel home page. Accessed January 20, 2015, http://www2a.cdc.gov/ilinet.

Centers for Disease Control and Prevention (2010b) Seasonal influenza (flu), weekly report: Influenza summary update. Accessed January 20, 2015, http://www.cdc.gov/flu/weekly/.

Chowell G, Echevarría-Zuno S, Viboud C, Simonsen L, Tamerius J, Miller MA, Borja-Aburto VH (2011) Characterizing the epidemiology of the 2009 influenza A/H1N1 pandemic in Mexico. *PLoS Medicine* 8(5):e1000436.

Edney A (2015) Flu vaccine effective in fewer than 1 in 4 in U.S. Bloomberg news. Accessed January 20, 2015, http://www.bloomberg.com/news/2015-01-15/flu-vaccine-effective-in-fewer-than-1-in-4-in-u-s-this-season.html.

Finkelstein SN, Hedberg KJ, Hopkins JA, Hashmi S, Larson RC (2011a) Vaccine availability in the United States during the 2009 H1N1 outbreak. *Amer. J. Disaster Medicine* 6(1):23–30.

Finkelstein SN, Prakash S, Nigmatulina K, McDevitt J, Larson RC (2011b) A home toolkit for primary prevention of influenza by individuals and families. *Disaster Medicine Public Health Preparedness* 5(4):266–271.

Flu.gov (2010) Pandemic flu history. Accessed January 20, 2015, http://www.flu.gov/pandemic/history/index.html.

Fu YC (2005) Measuring personal networks with daily contacts: A single-item survey question and the contact diary. *SOC Networks* 27:169–186.

Fu YC (2007) Contact diaries: Building archives of actual and comprehensive personal networks. *Field Methods* 19(2):194–217.
Hardin G (1968) The tragedy of the commons. *Science* 162(3859):1243–1248.
Hatchett RJ, Mecher CE, Lipsitch M (2007) Public health interventions and epidemic intensity during the 1918 influenza pandemic. *PNAS* 104(18):7582–7587.
Heesterbeek JAP (2002) A brief history of R_0 and a recipe for its calculation. *Acta Biotheoretica* 50(3):189–204.
Kaplan EH, O'Keefe E (1993) Let the needles do the talking! Evaluating the New Haven needle exchange. *Interfaces* 23(1):7–26.
Larson RC, Nigmatulina KR (2009) Engineering responses to pandemics. Rouse W, Cortese D, eds. *Engineering the System of Healthcare Delivery* (IOS Press BV, Amsterdam), 311–340.
Larson RC, Teytelman A (2012) Modeling the effects of H1N1 influenza vaccine distribution in the United States. *Value Health* 15(1): 158–166.
Lo JYC, Tsang THF, Leung Y-H, Yeung EYH, Wu T, Lim WWL (2005) Respiratory infections during SARS outbreak, Hong Kong, 2003. *Emerging Infectious Diseases* 11(111):1738–1741.
McKibbin W (2006) Global macroeconomic consequences of pandemic influenza. Lowy Institute for International Policy, Sydney. Accessed January 20, 2015, https://cama.crawford.anu.edu.au/pdf/working-papers/2006/262006.pdf.
McNamee D (2015) One-shot flu vaccine will be ready "within 5 years." *Medical News Today*. Accessed January 20, 2015, http://www.medicalnewstoday.com/articles/288082.php.
Mossong J, Hens N, Jit M, Beutels P, Auranen K, Mikolajczyk R, Massari M, et al. (2008) Social contacts and mixing patterns relevant to the spread of infectious diseases. *PLoS Medicine* 5(3):e74.
National Foundation for Infectious Diseases (2008) *Call to Action: Influenza Immunization Among Health Care Personnel* (National Foundation for Infectious Diseases, Bethesda, MD).
Nigmatulina K, Larson RC (2009) Living with influenza: Impacts of government imposed and voluntarily selected interventions. *Eur. J. Oper. Res.* 195(2):613–627.
Prachayangprecha S, Makkoch J, Payungporn S, Chieochansin T, Vuthitananchot C, Vuthitanachot V, Theamboonlers A, Poovorawan Y (2010) Serological analysis of human pandemic influenza (H1N1) in Thailand. *J. Health Population Nutrition* 28(6):537–544.
Rambhia KJ, Watson M, Sell TK, Waldhorn R, Toner E (2010) Mass vaccination for the 2009 H1N1 pandemic: Approaches, challenges, and recommendations. *Biosecurity Bioterrorism* 8(4):321–330.
Rouse W, Cortise DA (2010) *Engineering the System of Healthcare Delivery* (IOS Press, Amsterdam).
Sandman PM (2007) A severe pandemic is not overdue—It's not when but if. CIDRAP News and Perspective. Accessed January 20, 2015, http://www.cidrap.umn.edu/news-perspective/2007/02/severe-pandemic-not-overdue-its-not-when-if.
Stewart AM, Cox MA, O'Connor ME (2011) *Influenza Vaccination of the Health Care Workforce: A Literature Review School of Public Health and Health Services* (George Washington University, Washington, DC).
Teytelman A, Larson RC (2012) Modeling influenza progression within a continuous-attribute heterogeneous population. *Eur. J. Oper. Res.* 220(1):238–250.
Wein LM, Atkinson MP (2009) Assessing infection control measures for pandemic influenza. *Risk Analysis* 29:949–962.
Wu JT, Wein LM, Perelson AS (2005) Optimization of influenza vaccine selection. *Oper. Res.* 53(3):456–476.
Zhang D, Liu W, Yang P, Zhang Y, Li X, Germ KE, Tang S, Sun W, Wang Q (2013) Factors associated with household transmission of pandemic (H1N1) 2009 among self-quarantined patients in Beijing, China. *PLoS ONE* 8(10):e77873.