Proceedings of the 2010 Winter Simulation Conference
B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan, eds.

SIMULATION OF STRATEGIES FOR CONTAINING PANDEMIC INFLUENZA

Sigrún Andradóttir
Wenchi Chiu
David Goldsman
Mi Lim Lee
Kwok-Leung Tsui

Georgia Institute of Technology
H. Milton Stewart School of
Industrial and Systems Engineering
Atlanta, Georgia, U.S.A.

David N. Fisman

University of Toronto
Department of Epidemiology
Dalla Lana School of Public Health
Toronto, Ontario, CANADA

Beate Sander

University of Toronto
Department of Health Policy,
Management and Evaluation
Toronto, Ontario, CANADA

Azhar Nizam

Emory University
Department of Biostatistics
and Bioinformatics
Atlanta, Georgia, U.S.A.

ABSTRACT

We use a stochastic simulation model of pandemic influenza to investigate realistic intervention strategies that can be used in reaction to developing outbreaks. The model is constructed to represent a typical mid-sized North American city. Our model predicts average illness attack rates and economic costs for various intervention scenarios, e.g., in the case when low-coverage reactive vaccination and limited antiviral use are combined with minimally disruptive social distancing strategies, including short-term closure of individual schools. We find that such combination strategies can be substantially more effective than vaccination alone from epidemiological and economic standpoints.

1 INTRODUCTION

With modern advances in science and technology, simulation has been widely used in many research fields. Epidemic simulation models provide useful tools to increase our understanding of the dynamics and patterns of disease propagation, and to study and evaluate the potential impacts of various government policies and intervention strategies for pandemic diseases, including prophylactic use of antivirals and social distancing strategies such as school closure, quarantine, and isolation. These tools are especially timely in light of the recent H1N1 swine flu pandemic.

Similar to other simulation studies, epidemic simulations consist of three parts: identification of relevant inputs, model building, and output data analysis. The inputs typically include basic population data, the contact behavior of the individuals in the population, and disease transmission parameters. In the U.S., the population and their behavior are generated based on census bureau and demographic research, respectively, and the transmission parameters are based on epidemiologic research. Regarding model logic, the most-popular approaches are individual-based epidemiological compartment models, for example,

SIR (Susceptible-Infectious-Recovered) models. The simulation outputs are usually analyzed based on epidemiological and economic considerations. The epidemiological analysis often includes illness attack rates on various subsets of the populations, as well as the basic reproductive number (R_0). The quantity R_0 is defined as the expected number of secondary cases produced by a typical infected individual during its entire period of infectiousness in a completely susceptible population (Diekmann, Heesterbeek, and Metz 1990), and it plays a central role in mathematical estimation of disease spread (but is difficult to estimate via simulation). If one is also interested in an economic analysis, then all applicable social expenses, both direct and indirect, ought to be included in the mix.

We built a simulation model of pandemic influenza to investigate realistic strategies that can be used in reaction to developing outbreaks. The simulator, which is similar to models developed by Longini et al. (2004, 2005) and Aleman, Wibisono, and Schwartz (2009), is programmed in C++ and runs on desktop platforms. Our model is calibrated to documented illness attack rates and basic reproductive number estimates of 2009 H1N1 (Swine flu), and is constructed to represent a typical mid-sized North American city. We consider both medical and non-medical disease mitigation strategies in the model. Our model predicts average illness attack rates and economic costs under various intervention scenarios. We find that certain "combination" strategies can be substantially more effective than vaccination alone from epidemiological and economic standpoints.

The remainder of the paper is organized as follows. In Section 2, we present the model structure. Section 3 concerns the model's input parameters, including some remarks on the issue of calibration. In Section 4, we study a number of reasonable intervention strategies. Section 5 deals with output analysis, both epidemiological and economic, and Section 6 offers some conclusions. Parts of this paper are taken from Andradóttir et al. (2010).

2 MODEL STRUCTURE

This section describes the general model structure underlying the simulation. First of all, we discuss the transmission route of the disease. We use a well-known methodology for modeling airborne infectious disease in a small area, namely, the SIR ("susceptible–infectious–removed'') compartmental model. The structure of the SIR model is depicted in Figure 1. Every individual in the population has four possible states: susceptible, infected but not infectious, infectious, and removed (either dead, or recovered with immunity).

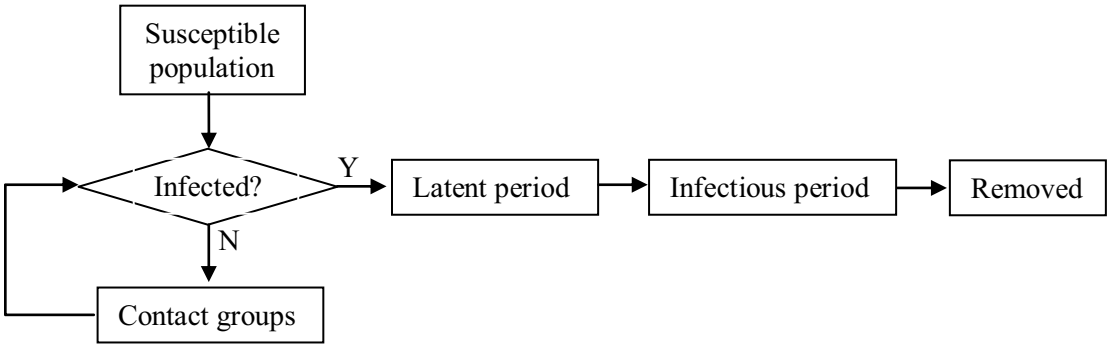

Figure 1: Flowchart of the epidemiological SIR model

An outbreak is initialized by randomly generating infected individuals at the beginning of day 1, with all other individuals considered susceptible. Susceptible people have the opportunity, each day, to become infected in their contact groups. Each person's contact groups are simply the set of groups with whom the person may interact during a particular day. The contact groups we considered include household, neighborhood, community, school (daycare, playground, elementary, middle, and high), and workgroup. In our model, everyone has contacts with people in their household, neighborhood, and commu-

nity. Children of ages 0–4 may also visit a daycare center or a playground (50% go to daycare and 50% go to a playground in our model). Children of ages 5–18 attend schools in the same community and, of course, contact other students in the same school. In addition, some children aged 16–18 may be employed, and thus they stay in a workgroup instead of school. Workers contact all of the persons in their workgroups. Employment rates are based on available employment statistics as well as census statistics on the geographical distribution of workers; details are given in Andradóttir et al. (2010).

We set up our community structure based on the framework given in Longini et al. (2004). Each community consists of approximately 2000 people who live in the same census tract; population age and household type distributions are based on Year 2000 Census information. Each community is broken down into 4 neighborhoods, with a total of 2 elementary schools, 1 middle school, and 1 high school. We formed workgroups of size 20 to represent the typical number of people with whom a person has close contact during the day.

The simulator tracks the infection status of each person in the population during each time period (day), and the infection and illness status of each person are recorded, as are the illness attack rates overall and in various age groups. The system state gets updated at fixed time intervals (at the end of each day). The daily probability of infection for each susceptible person is determined by the number of infectious contacts in his contact groups, and on the per-contact probability of transmission for each type of contact. For example, the probability of infection on a particular day for a susceptible child who attends daycare is

$$\begin{aligned} 1 - [&\Pr(\text{child is not infected in the household}) \\ &\times \Pr(\text{child is not infected in the neighborhood}) \\ &\times \Pr(\text{child is not infected in the community}) \\ &\times \Pr(\text{child is not infected at the daycare center})]. \end{aligned}$$

Within each contact group, the probability of infection of a susceptible individual depends on the number of infectious individuals in the group. For example, suppose that X children and Y adults in a household are infectious on a particular day. Then the probability of a susceptible household member being infected in that household on that day is (assuming independence among infected household members):

$$1 - [\Pr(\text{not infected by a child in the household})^X \times \Pr(\text{not infected by an adult in the household})^Y].$$

The number of infectious people in the contact groups (X and Y) are random variables that are updated at the beginning of each day.

If the susceptible person is not infected, he continues with his routine contacts until the next check time. People infected with disease first pass through a latent period (the period between being infected and becoming infectious). Due to the lack of historical epidemiological data, most researchers assume that the latent period is the same as the incubation period (the period between being infected and the appearance of symptoms). After the latent period, an infected person undergoes an infectious period. During the infectious period, the infected person may or may not develop influenza symptoms. Based on the work of Weycker et al. (2005) and the references therein, we set the probability that an infected person will exhibit symptoms to a value of 0.67; the determination of such probabilities is an ongoing research problem in the field. After the infectious period, people enter the removed state, where they are either recovered (with immunity for this particular strain of influenza), hospitalized, or dead.

3 INPUT CONSIDERATIONS

From our model description in Section 2, we recall that the inputs to the simulation include various population characteristics (e.g., population age distribution and contact behaviors), transmission probabilities, and durations of the latent and infectious periods. Our population is generated using a combination of census bureau data and resident commuting data. The population's contact structure is similar to that discussed in Longini et al. (2004). The transmission probabilities are based on two sources. First, for

transmissions within households, we again use results from Longini et al. (2004). For schools, workplaces, neighborhoods, and communities, we calibrated transmission probabilities (simulation inputs) so that the realized illness attack rates and R_0 values (simulation outputs) matched those from Tuite et al. (2010) for the 2009 H1N1 pandemic; specifically, $R_0 = 1.4$. Concerning the durations of the latent and infectious periods, we simply used the empirical cumulative distributions from Longini et al. (2004, 2005); see Table 1, where the lengths of the latent and infectious periods are assumed to follow empirical probability distributions with means of 1.9 and 4.1 days, respectively.

Table 1: Empirical distributions of latent and infectious periods for flu (Longini et al. 2004, 2005)

Latent period		Infectious period	
Duration (days)	Cumulative probability	Duration (days)	Cumulative probability
1	0.3	3	0.3
2	0.8	4	0.7
3	1.0	5	0.9
		6	1.0

The process of calibrating transmission probabilities to realized illness attack rates and R_0 turns out to be a tedious iterative manual exercise, illustrated in Figure 2. The idea is to input reasonable transmission probabilities for various age groups; see if the output matches the observed R_0 and age-specific attack rates (in this case, from Tuite et al. 2010); and then iteratively tweak the inputs to obtain a better match. The difficulty of calibration is that the number of inputs (each transmission rate for each age/contact group) is daunting, and the inputs tend to be highly correlated. The problem of conducting efficient calibration is one of ongoing research.

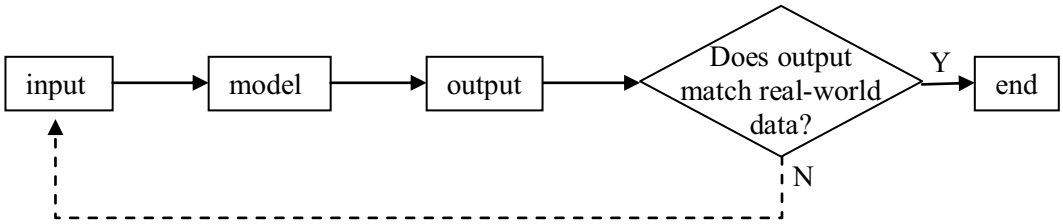

Figure 2: The calibration process

In our model, the calibrated attack rates were 29.5% for children ages 0–4, 55.9% for children ages 5–13, 40.8% for adults ages 19–52, 14.3% for adults ages 53–59, and 11.0% for seniors over 60 (which reflects greater immunity of older individuals to H1N1). The calibrated attack rates and R_0 value are very close to the observed quantities. The resulting transmission probabilities — the ones actually used in our simulation — are shown in Table 2.

4 INTERVENTIONS

Our simulation models the baseline (no intervention) case along with various intervention combinations. The interventions include vaccination, antiviral treatment and household prophylaxis, and school closure and general social distancing, which are discussed in Sections 4.1, 4.2, and 4.3, respectively. Interventions are triggered when the overall illness attack rate reaches 0.01%.

4.1 Vaccination

We model a single-dose vaccine scenario based on the finding of Greenberg et al. (2009) that a single dose elicited a robust immune response in a large percentage of adults. We model both pre-vaccination as

well as reactive strategies, with reactive vaccination programs beginning immediately, 30 days, or 60 days after the outbreak starts. The delays model disruptions in the vaccine production and distribution supply chain. Our model also implements two possible vaccine stockpile levels (covering 35% and 70% of the population) and two sets of vaccine efficacies. Under the low-efficacy scenario considered here, the vaccine efficacy against susceptibility to infection (VE_S) is 0.3 (i.e., the probability that a person receiving the vaccine will become immune is 0.3), and similarly the vaccine efficacy against infectiousness (VE_I) is 0.2. We also consider the high-efficacy case, with $VE_S = 0.4$ and $VE_I = 0.5$. These values are consistent with estimates of influenza vaccine efficacy used by Basta et al. (2009) in their analysis of the impact of vaccinating schoolchildren against pandemic and seasonal influenza. Each day, we randomly select unvaccinated and non-ill people to receive vaccine based on the availability of doses. Vaccine protection builds over time, with half of the vaccine's efficacy realized upon vaccination, and full protection after two weeks.

Table 2: Per-contact influenza infection transmission probabilities within contact groups

Contact Group	Transmission Probability
Household	
Child to Child	0.8
Child to Adult	0.3
Adult to Child	0.3
Adult to Adult	0.4
Community[1]	
Pre-schooler	0.000005
School child	0.000005
Adult (ages 19–52)	0.000075
Adult (ages 53+)	0.000055
Daycares/Playgroups[1]	
Daycares	0.028
Playgroups	0.018
Schools[1]	
Elementary schools	0.012
Middle schools	0.011
High schools	0.010
Workgroups	0.010

[1]Probability that a susceptible person in the group is infected through contact with an infectious person in the group.

4.2 Antiviral Treatment and Household Prophylaxis

In our simulation, we modeled two types of antiviral use: one for purposes of treatment, and one for household prophylaxis (prevention of additional household infections). Antivirals are assumed to be stockpiled before the outbreak. We also assumed that antiviral courses are available for 10% of the population and that they are distributed to infected individuals and their household members until the supply is exhausted. Further, antivirals are given in five-day courses (which is common practice), and 1% of those people taking antivirals do not complete their entire courses due to side effects, inconvenience, etc. In addition, we used an antiviral efficacy against susceptibility (AVE_S) of 0.3 and an efficacy against infectiousness (AVE_I) of 0.7.

4.3 School Closure and Social Distancing

Our model combines school closure (closing schools for several days) and social distancing (people automatically reduce contacts in the workplace, neighborhood, and community after becoming aware of the outbreak) in the belief that school closure will immediately trigger social distancing, so that the two interventions will happen simultaneously. We assume a school closure duration of five days to be consistent with our goal of modeling limited and practical social distancing; and since the average length of the infectious period in our model is 4.1 days, a shorter closure period would probably not make sense. School closure is triggered by five children appearing with symptoms on any given day; the choice of a trigger point of five was arbitrary, but deemed reasonable. When school closure occurs, we also model a reduction in workplace and general community contacts of 20%.

5 OUTPUT ANALYSIS

The outputs from our epidemic simulation are used to analyze the effectiveness of various interventions. We provide both epidemiological and economic output analysis. The epidemiological analysis includes illness attack rates. The economic analysis includes all of the costs incurred during the outbreak, including both direct and indirect expenses.

We first discuss the epidemiological analysis. Our calibrated simulation shows that with no interventions, the average overall illness attack rate is 34.1%, which is close to the observed data (thus helping to validate the simulation). The strategy of pre-vaccination (V) of 35% of the population with a low-efficacy vaccine reduces the average overall illness attack rate to 26.1%. Pre-vaccination of 35% of the population combined with antiviral treatment and household antiviral prophylaxis (V+A) reduces the average overall illness attack rate to 19.3%. If school closure and social distancing is also added to the mix of interventions (V+A+S), then the average overall illness attack rate can be reduced to 1.0%. Pre-vaccination of 70% with a low-efficacy vaccine reduces the overall illness attack rate to 0.2% if combined with household antiviral prophylaxis and school closure and social distancing. Figure 3 depicts the average overall attack rates from our simulations under the various intervention strategies outlined above. The simulations also take into account various supply chain delays in delivering vaccination (pre-vaccination, reactive delay of 0 days, 30-day delay, and 60-day delay); we see that attack rates increase as delays increase, sometimes substantially. The results in Figure 3(a) incorporate vaccine dosages for 35% population coverage, and (b) for 70% population coverage. Obviously, the attack rates significantly improve as vaccine coverage goes from 35% to 70%. The analogous results for high-efficacy vaccine (not described here) are similar; we note that the high-efficacy vaccine can reduce most of the attack rates and economic costs.

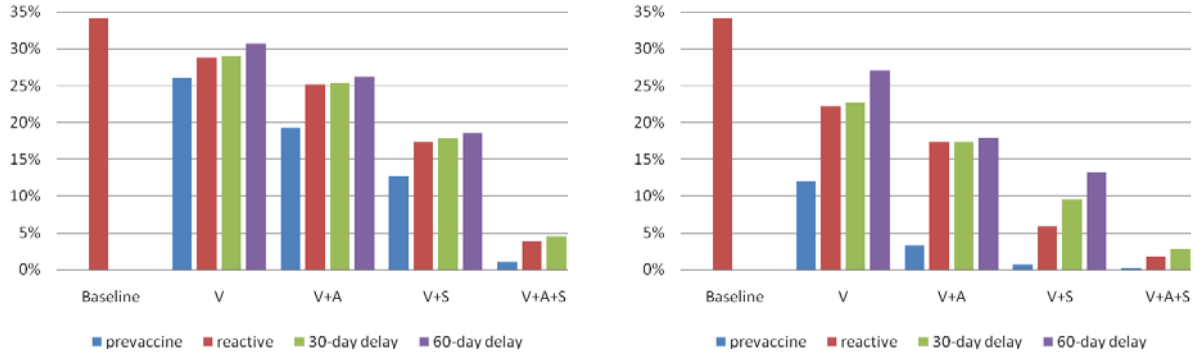

(a) 35% vaccine coverage

(b) 70% vaccine coverage

Figure 3. Average overall illness attack rates under different intervention strategies

Now we discuss the economic analysis, where we use methods described by Meltzer, Cox, and Fukuda (1999) to quantify most medical and work loss costs (see also Medlock and Galvani 2009). The economic analysis is based on three tables, all of which are described in full detail in Andradóttir et al. (2010): (i) the proportions of cases at high risk for complications (by age group); (ii) the rates of outpatient visits, hospitalizations, and death for people having complications; and (iii) frequency and costs associated with outpatient visits, hospitalizations, and deaths, e.g., the average number of visits per case, average copayment per outpatient visit, average copayment per prescription, average value of total days of work lost, etc. In the latter table, all the costs have been adjusted for inflation using 2008 consumer price and medical price indexes. In our analysis, we chose the "low" rate estimates presented in Meltzer, Cox, and Fukuda (1999), which we believe to be most consistent with the relatively low R_0 (1.4) for our model.

Our simulation shows that with no intervention, the total cost during the outbreak will be $81.1 million. Pre-vaccination (V) of 35% of the population with a low-efficacy vaccine reduces the cost to $71.1 million. Pre-vaccination of 35% of the population combined with antiviral treatment and household antiviral prophylaxis (A) reduces the cost to $56.4 million. If school closure and social distancing (S) are also added, the cost can be reduced to $15.9 million. Pre-vaccination of 70% with a low-efficacy vaccine, which will increase the vaccine cost, can reduce economic cost to $21.3 million if combined with household antiviral prophylaxis and school closure and social distancing. See Figure 4, where the y-axis is in units of US million dollars.

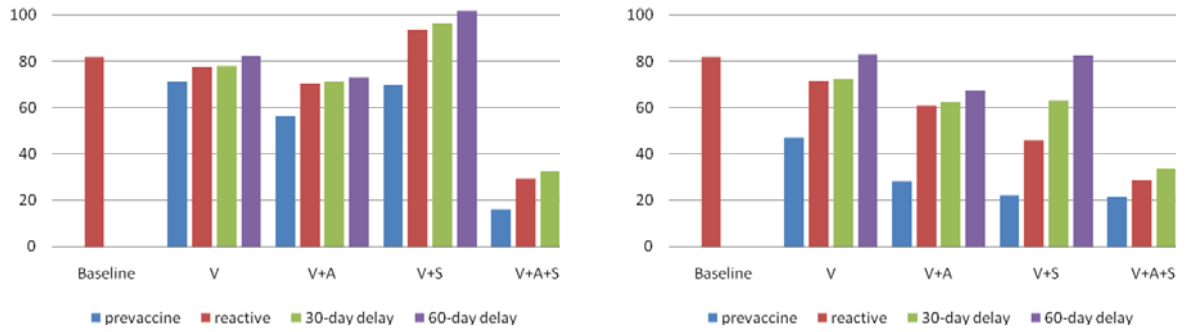

(a) 35% vaccine coverage

(b) 70% vaccine coverage

Figure 4. Economic analysis under different policies

Finally, we examine the epidemiological and economic analyses together. Figure 3 shows that adding intervention options, as in the sequence V, V+A, V+S, V+A+S, decreases the attack rates. However, from Figure 4(a), we see that for the case of 35% vaccine coverage, the economic costs for certain intervention strategies may end up being higher than the baseline (no intervention) strategy. For example, for the V+S intervention strategy, the savings associated with the decreased attack rates are not enough to balance the costs of vaccine and work loss. Similar examples can be observed for the 70% vaccine coverage case.

6 CONCLUSIONS

In this paper, we use a stochastic simulation model of pandemic influenza to investigate realistic intervention strategies that can be used in reaction to developing outbreaks. Our simulation model can represent a typical mid-sized North American city and predicts average illness attack rates and economic costs under various intervention scenarios; specifically, low-coverage reactive vaccination and limited antiviral use are combined with minimally disruptive social distancing strategies, including short-term closure of indi-

vidual schools. We find that all the interventions can decrease attack rates. Of course, from an economic perspective, the strategies may result in higher or lower costs than doing nothing (baseline).

REFERENCES

Aleman, D. M., T. G. Wibisono, and B. Schwartz. 2009. Accounting for individual behaviors in a pandemic disease spread model. In *Proceedings of the 2009 Winter Simulation Conference*, ed. M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, and R. G. Ingalls, 1977–1985. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Andradóttir, S., W. Chiu, D. Goldsman, M. L. Lee, K.-L. Tsui, B. Sander, D. N. Fisman, and A. Nizam. 2010. Reactive strategies for containing developing outbreaks of pandemic influenza. Technical report, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia.

Basta, N. E., D. L. Chao, M. E. Halloran, L. Matrajt, and I. M. Longini. 2009. Strategies for pandemic and seasonal influenza vaccination of schoolchildren in the United States. *American Journal of Epidemiology* 170:679–686.

Diekmann, O., J. A. Heesterbeek, and J. A. Metz. 1990. On the definition and the computation of the basic reproduction ratio R_0 in models for infectious disease in heterogeneous populations. *J. Math. Biol.* 28(4):365–382.

Greenberg, M. E., M. H. Lai, G. F. Hartel, C. H. Wichems, C. Gittleson, J. Bennet, G. Dawson, W. Hu, C. Leggio, D. Washington, and R. L. Basser. 2009. Response after one dose of a monovalent influenza A (H1N1) 2009 vaccine — preliminary report. *New England Journal of Medicine* 361.

Longini, I. M., M. E. Halloran, A. Nizam, and Y. Yang. 2004. Containing pandemic influenza with antiviral agents. *American Journal of Epidemiology* 159:623–633.

Longini, I. M., A. Nizam, S. Xu, K. Ungchusak, W. Hanshaoworakul, D. A. T. Cummings, and M. E. Halloran. 2005. Containing pandemic influenza at the source. *Science* 309:1083–1087.

Medlock, J., and A. P. Galvani. 2009. Optimizing influenza vaccine distribution. *Science* 325:1705–1708.

Meltzer, M. I., N. J. Cox, and K. Fukuda. 1999. The economic impact of pandemic influenza in the United States: Implications for setting priorities for intervention. *Emerging Infectious Diseases* 5:659–671.

Tuite, R. A., A. L. Greer, M. Whelan, A. Winter, B. Lee, P. Yan, J. Wu, S. Moghadas, D. Buckeridge, B. Pourbohloul, and D. N. Fisman. 2010. Estimated epidemiologic parameters and morbidity associated with pandemic H1N1 influenza. *Canadian Medical Association Journal* 182(2):131–136.

Weycker, D., J. Edelsberg, M. E. Halloran, I. M. Longini, A. Nizam, V. Ciuryla, and G. Oster. 2005. Population-wide benefits of routine vaccination of children against influenza. *Vaccine* 23:1284–1293.

AUTHOR BIOGRAPHIES

SIGRÚN ANDRADÓTTIR is a Professor in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. Prior to joining Georgia Tech, she was an Assistant Professor and later an Associate Professor in the Departments of Industrial Engineering, Mathematics, and Computer Sciences at the University of Wisconsin–Madison. She received her Ph.D. in Operations Research from Stanford University in 1990. Her research interests include simulation, applied probability, and stochastic optimization. She is a member of INFORMS and served as Editor of the *Proceedings of the 1997 Winter Simulation Conference*. She was the Simulation Area Editor of *Operations Research Letters* from 2002 to 2008, and has served as Associate Editor for various journals. Her e-mail address is <sa@gatech.edu>, and her web page is <http://www.isye.gatech.edu/faculty/sa>.

WENCHI CHIU is a post-doctoral researcher in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. Her research interests include applied statistics, si-

mulation analysis, and healthcare simulation. Her email address is <wenchichiu@gmail.com>, and her web page is <http://www.prism.gatech.edu/~gtg979v/>.

DAVID GOLDSMAN is a Professor in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. His research interests include simulation output analysis, ranking and selection, and healthcare simulation. He is an active participant in the Winter Simulation Conference, having been Program Chair in 1995, and having served on the WSC Board of Directors from 2002 to 2010. His e-mail address is <sman@gatech.edu>, and his web page is <www.isye.gatech.edu/~sman>.

MI LIM LEE is a third-year Ph.D. student in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology. Her research focus has been on healthcare simulation and biosurveillance. Her email address is <mlee79@gatech.edu>.

KWOK-LEUNG TSUI is a Professor in the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology and a Chaired Professor of Industrial Engineering at the City University of Hong Kong. He received his Ph.D. in Statistics from the University of Wisconsin–Madison in 1986. Dr. Tsui was a recipient of the NSF Young Investigator Award in 1992. He was the (elected) President and Vice President of the American Statistical Association Atlanta Chapter in 1992–1993; the Chair of the INFORMS Section in Quality, Statistics, and Reliability in 2000; and the Founding Chair of the Section in Data Mining in 2004. He is a Fellow of the American Statistical Association, a U.S. representative in the ISO Technical Committee on Statistical Methods, and a Department Editor of *IIE Transactions*. His current research interests include health informatics, data mining and surveillance in healthcare and public health, calibration and validation of computer models, bioinformatics, process control and monitoring, and robust design and the Taguchi method. His e-mail address is <ktsui@isye.gatech.edu>, and his web page is <www.isye.gatech.edu/~ktsui>.

DAVID FISMAN is an Associate Professor of Infectious Disease Epidemiology in the Dalla Lana School of Public Health at the University of Toronto. He is interested in the epidemiology and dynamics of respiratory infections and sexually transmitted infections, and the intersection between environmental change and infectious disease risk. His email address is <david.fisman@utoronto.ca>.

BEATE SANDER is a health economist with the Ontario Agency of Health Protection and Promotion, a collaborator with the Toronto Health Economics and Technology Assessment Collaborative, and a Ph.D. candidate at the Department of Health Policy, Management and Evaluation, University of Toronto. She is interested in answering public health policy questions by using models to assess the effectiveness and cost-effectiveness of healthcare programs. She has extensive experience applying advanced methods to decision problems in the prevention and treatment of influenza, as well as to other clinical indications, technologies, and programs such as cervical cancer screening, osteoporosis, and surgical techniques. Her recent projects include economic evaluations of Ontario's Universal Influenza Immunization Program and of influenza pandemic mitigation strategies in Canada and the U.S. Her email address is <beate.sander@oahpp.ca>.

AZHAR NIZAM is a faculty member in the Rollins School of Public Health of Emory University. His research interests include stochastic simulation of infectious disease outbreaks, analysis of clinical trials data, and statistics education. He has spent much of the last decade developing and reporting results of simulations of seasonal and pandemic influenza, avian influenza, cholera, and smallpox. His email address is <anizam@emory.edu>.