

A LARGE SIMULATION EXPERIMENT TO TEST INFLUENZA PANDEMIC BEHAVIOR

Michael F. Beeler
Dionne M. Aleman
Michael W. Carter

Department of Mechanical and Industrial Engineering
University of Toronto
5 King's College Road
Toronto, M5S 3G8, CANADA

ABSTRACT

The effectiveness of mass vaccination and voluntary self-quarantine to mitigate pandemic influenza is tested in a large agent-based simulation. The characteristics of the pandemic—infectiousness, days contagious, and risk of death—are varied systematically along with the mitigation efforts in a five-factor designed experiment. A total of 243 distinct pandemic scenarios are tested. A range of two-way and three-way interaction effects are found that show significant non-linearities and contingencies in pandemic behavior and intervention effectiveness.

1 INTRODUCTION

Agent-based simulations of pandemic influenza have grown in popularity as a tool for predicting the course of possible influenza pandemics and evaluating mitigation strategies (Aleman et al. 2011; Haber et al. 2007; Los Alamos National Laboratory 2006; Das et al. 2008). While these highly detailed simulations offer the ability to manipulate many variables, they are often computationally expensive, especially when modeling large urban areas. For example, a pandemic simulation called EpiSimS developed by the Los Alamos National Laboratories in the United States, the most detailed developed to date, requires 45 minutes to initialize its population and 12 hours of run time to generate results for 160 simulated days (2006). Another simulation developed by Das et al. (2008) requires 50 minutes to simulate one day for a populations of 1.1 million on one processor. Consequently, most agent-based pandemic simulation studies have not yet explored the full diversity of scenarios that they are capable of testing, and were designed to test.

There has been, however, some progress in this area. Das et al. have conducted a multi-factorial simulation experiment on a population of 100,000 (2008), and some limited scenario testing is presented in a agent-based vaccine cost-effectiveness study (Durbin et al. 2011), and in non-agent-based cost-effectiveness studies (Khazeni et al. 2009; Sander et al. 2010). We are unaware of any agent-based simulation experiments on a large scale, in terms of population size and number of scenarios tested, conducted to date.

To bridge this gap, we conducted a five variable full-factorial experiment, testing 243 distinct scenarios, with six replications, using an agent-based pandemic simulation program. The program is described in detail in (Aleman et al. 2011), and was used for one-way scenario testing in (Beeler et al. 2011). The objective of the study was not to arrive at specific point estimates for the cost-effectiveness or impact of an intervention, but to test hypotheses about the presence and general magnitude of interaction effects, specifically of combining two policies, and the circumstances that tend to accentuate or diminish their relative effectiveness. The factors tested were vaccination levels, rates of self-isolation, baseline infectiousness, average days contagious, and the risk of death. Each factor was tested at three levels. The results reveal

important interactions between the characteristics of the pandemic and the effectiveness of mitigation measures.

2 AGENT-BASED PANDEMIC INFLUENZA MODEL

The simulation program works by assigning 5 million individuals to households, census tracts, age groups, and workplaces based on 2006 census data for the Greater Toronto Area (Ontario, Canada). Agents have contact with members of their household, workplaces, and community. Contact rates were based on a variety of sources discussed in (Aleman et al. 2011). The average probability of disease transmission per minute of contact between susceptible and infectious persons is varied by age group and was estimated by Haber et al. (2007) and Valle et al. (2007), and adapted to the context considered in our model by epidemiologists from the Ontario Agency for Health Protection and Promotion (OAHPP).

The simulation initializes with seven infectious individuals, and computes the probability that each person having contact with an infected individual on a given day will become infected. The simulation runs for 100 days and outputs the number of infections and deaths. The duration of infection and probability of death depends on the person's age group. Those who do not die from influenza are considered recovered, and are not able to contract the virus again or infect others. The duration of infectiousness, susceptibility, and behavior patterns also vary by age group. For simplicity, it was assumed that no members of the population had immunity prior to the pandemic's onset.

3 DESIGNED EXPERIMENT

The designed experiment tests two mitigation strategies and three factors that represent variability in a pandemic's intrinsic characteristics. The mitigation responses tested are mass vaccination and levels of self-quarantine by people who become symptomatically infected. The vaccination schemes tested included the following: no vaccination; vaccination of 30% of the population in equal increments over 60 days; and vaccination of 60% of the population in equal increments over 60 days. For the purposes of testing and simplicity, the vaccination schemes were assumed to begin on the first day of the pandemic. In reality, it may take a couple months before vaccines are developed and approved for distribution from the onset of the first global cases of a pandemic flu strain, as was the case with pH1N1 in 2009. When infected individuals self-quarantine, they may still infect other members of their household, but are assumed to have no further contact with the community. The levels of voluntary self-quarantine tested were 0%, 30%, and 60%.

The three characteristics of the pandemic that were varied are 1) the average number of days an infected person is contagious; 2) the degree of infectiousness; and 3) and the risk of death from infection. Infected individuals were assumed to become contagious the day after becoming infected. In practice, incubation periods may vary by person, and could be longer depending on the virus strain. The baseline number of days contagious was set to range uniformly between four and five days in persons under 20, four and eight days in persons under 40, and three to 10 days in persons 40 and over. The two other factor levels for days contagious were obtained by shifting the upper and lower bounds of the possible days contagious down by two days in one case, and up by two days in the other case, for all age groups.

The infectiousness factor represents the risk of infection per unit of contact time between infectious and susceptible persons of particular age groups. The baseline infectiousness parameters were drawn from (Haber et al. 2007; Valle et al. 2007) and used in the earlier published studies involving this simulation program (Aleman et al. 2011; Beeler et al. 2011). The three levels of infectiousness were a 50% decrease from the baseline, no change, and a 50% increase from the baseline.

The simulation evaluates whether an infected person will recover, remain infected, or die, at the end of each day. The probability of death on a given day was treated as a constant for simplicity. The three levels for this factor are: 0.001, 0.005, and 0.025. A wide range was used because the risk of death from influenza can vary substantially from strain to strain. The risk of death was varied independently of other factors and

was not programmed to give rise to special behavior; in practice, it may be that rates of self-quarantine would increase for more deadly virus strains.

The response variable used is the total number of infections occurring within the first 100 days of the pandemic

4 RESULTS

The three virus-property factors were given low, medium, and high levels. Subscripts 1 and 2 next to the parameter symbols in Table 1 denote low and high levels for these factors, respectively. The three medium levels were used in the baseline scenario, and do not have subscripts since their presence is implied by omission of an adjustment parameter. For the two human response factors, no vaccination and no self-quarantine were used to generate the baseline scenario. Subscripts 1 and 2 represent the 30% and 60% levels of vaccination and self-quarantine. With 243 scenarios and six replications, there were sufficient degrees of freedom to test for three-way interaction effects. All parameters presented in the table below represent deviations in the number of infections from the baseline scenario. The estimate for the baseline scenario was 40,437 infections. Interaction parameters are represented with two contiguous greek letters corresponding to the ones used for the main effects of the same factors.

In the initial data analysis, several models were fit to the data, including a model with main effects only, one with all two-way interactions, and one with all-three way interactions. In all three models, the risk of death did not have a major effect on total infections, and so new models were fit to the data that dropped this factor and its interactions. The fact that changes in the risk of death, at the levels tested, did not impact infections is instructive, because while it would theoretically impact the average amount of time people spend in a contagious state, we may now infer that such effects could be negligible even when the risk of dying reaches 2.5% per day, or over 10% per case on average.

The simplified three-way interaction model provided a much higher adjusted R-squared value (0.8762) than the two-way interaction model (0.7601), and the main effects model (0.489). All models had very low p -values ($p < 2.2\text{e-}16$). The three-way model provided the best fit to the data, and its parameters are presented in Table 1. A model with statistically significant interaction effects means that the effect of a change in a parameter value is contingent on the values of other parameters with which it interacts. The presence of interactions cannot be detected when variables are varied one at a time, and gives further motivation to using large multivariable simulation experiments to study the potential dynamics of pandemic influenza.

It is important to note that individual interaction parameters and main effects often cannot be interpreted meaningfully by themselves. However, in this particular case, the omission of the baseline factor levels from the parameter notation means that the main effects, and the main effects plus just the two-way interactions, actually constitute fully specified scenarios because all factors not part of the interaction term or main effect are implicitly set to their baseline values.

Because of the number of scenarios tested, a detailed text discussion of each scenario is not possible. Instead, a subset is presented in detail, which should facilitate interpretation of the scenarios not discussed in this manuscript for readers who are less familiar with three-way interaction models.

4.1 Parameter Interpretation

Vaccinating 30% and 60% of the population over 60 days reduced the number of infections occurring by 15,990 and 24,380 infections respectively, assuming no other changes in the baseline scenario. Likewise, the 30% and 60% self-quarantine rates reduced infections by 7,252 and 11,599. The effect of a 60% self-quarantine rate in combination with 30% or 60% vaccination resulted in 7,051 and 9,297 fewer prevented infections than the sum of the interventions' individual impact.

The effect of increasing infectiousness by 50%, all else held at baseline, is an increase of 70,360 infections. The higher infection rate also amplifies the magnitude of the effects of the two interventions. Vaccinating

Table 1: Results from five-factor designed experiment—statistically significant parameters only.

Parameter	Symbol	Estimate	p-value
Baseline	-	40,437	<0.001
Vaccination, 30%	α_1	-15,990	<0.001
Vaccination, 60%	α_2	-24,380	<0.001
Self-quarantine, 30%	β_1	-7,252	0.016
Self-quarantine, 60%	β_2	-11,599	<0.001
Days contagious -2	η_1	-12,517	<0.001
Days contagious +2	η_2	14,337	<0.001
Infectiousness -50%	λ_1	-25,502	<0.001
Infectiousness +50%	λ_2	70,360	<0.001
Two-way interactions			
Vaccination and Self-quarantine	$\alpha_2\beta_1$	5,778	0.081
	$\alpha_1\beta_2$	7,051	0.033
	$\alpha_2\beta_2$	9,297	0.005
Vaccination and Days contagious	$\alpha_2\eta_1$	7,586	0.022
	$\alpha_1\eta_2$	-5,730	0.083
	$\alpha_2\eta_2$	-8,791	0.0079
Vaccination and Infectiousness	$\alpha_1\lambda_1$	11,977	0.0051
	$\alpha_2\lambda_1$	17,233	<0.001
	$\alpha_1\lambda_2$	-50,251	<0.001
	$\alpha_2\lambda_2$	-65,348	<0.001
Self-quarantine and Infectiousness	$\beta_2\lambda_1$	10,340	0.016
	$\beta_1\lambda_2$	-41,902	<0.001
	$\beta_2\lambda_2$	-55,920	<0.001
Days contagious and Infectiousness	$\eta_1\lambda_1$	10,292	0.016
	$\eta_2\lambda_1$	-11,554	0.0069
	$\eta_1\lambda_2$	-26,172	<0.001
	$\eta_2\lambda_2$	101,460	<0.001
Three-way interactions			
Self-quarantine, Days contagious and Infectiousness	$\beta_2\eta_1\lambda_2$	12,127	0.0096
	$\beta_1\eta_2\lambda_2$	-46,950	<0.001
	$\beta_2\eta_2\lambda_2$	-56,964	<0.001
Vaccination, Days contagious and Infectiousness	$\alpha_1\eta_1\lambda_2$	10,980	0.019
	$\alpha_2\eta_1\lambda_2$	15,659	<0.001
	$\alpha_1\eta_2\lambda_2$	-52,714	<0.001
	$\alpha_2\eta_2\lambda_2$	-61,820	<0.001
Vaccination, Self-quarantine and Infectiousness	$\alpha_2\beta_2\lambda_1$	-8,356	0.074
	$\alpha_1\beta_1\lambda_2$	43,261	<0.001
	$\alpha_2\beta_1\lambda_2$	51,175	<0.001
	$\alpha_1\beta_2\lambda_2$	53,802	<0.001
	$\alpha_2\beta_2\lambda_2$	64,735	<0.001

60% of the population prevented 65,348 more infections than the number prevented in the baseline, because there were more infections to be prevented. The same holds true for the effect of 60% self-quarantine rates, which prevented 41,902 more infections than when infectiousness was unadjusted. However, as was the case under baseline infectiousness, the effect of implementing both policies simultaneously is far less than the sum of their individual impact. To be exact, the gap between the sum and the outcome from simultaneous intervention is 64,735 wider than under baseline infectiousness. Given 60% vaccination coverage, the marginal effect of the 60% self-quarantine rate, which can be calculated by adding β_2 , $\alpha_2\beta_2$, $\beta_2\lambda_2$, and $\alpha_2\beta_2\lambda_2$, is not statistically significant. What this means is that an intervention that may have a large impact by itself, and a lesser but still important added impact in conjunction with another intervention, may have little or no extra effect at all when the epidemiological characteristics of the pandemic change.

Whereas the interaction between the two mitigation actions resulted in reduced total effects, the interaction between days contagious and infectiousness shows a synergy. Whereas increasing average days contagious by 2 increases infections by 14,337 from the baseline, and increasing infectiousness by 50% pushes infections up 70,360, an additional 101,460 infections on top of the 14,337 and 70,360 occur when these two factors are applied simultaneously. This steep spike in additional infections, however, is tempered by 56,964 when the self-quarantine rate is set to 60%, as shown by the three-way interaction term.

5 DISCUSSION

The goal of this paper is not to establish point estimates for the effectiveness of a particular pandemic mitigation strategy, or to predict how severe a future pandemic may be. Those two unknowns have been subject to much investigation, and are to some degree unknowable because of the uncertainties and contingencies that characterize influenza pandemics. Rather, this paper takes a careful look at possible interaction effects between pandemic characteristics and mitigation strategies in order to gain insights of a more generalizable nature.

At a high level, the results show a number of things that are consequential to future pandemic modeling research. The presence of a large number of strong interaction effects supports this study's hypothesis that the effectiveness of a public health intervention may be contingent on the other interventions being implemented, and on the epidemiological properties of the pandemic. Studies that test interventions one-at-a-time, or that only conduct one-way sensitivity analysis, risk overlooking these interactions and mischaracterizing the impact we can expect pandemic interventions to have. Furthermore, scaling up interventions seems to produce diminishing marginal returns. In addition to the diminishing returns from combining strategies discussed in the previous section, the impact of 30% self-quarantine and 30% vaccination coverage, though half the other level tested, yielded in many cases 75-85% as much infection reduction as the 60% level. From an economic point of view, the marginal benefit of moving rates of vaccination or self-quarantine to 60% from 30% may be fairly small, especially when considering the higher marginal cost of going from a 30% vaccination and self-quarantine rate to a 60% rate, as the most eager or conscientious citizens would likely comprise the first 30%.

6 CONCLUSION

The current trend in research on influenza pandemic control has been to study parameter changes or interventions one-at-a-time, with many studies focusing on a single intervention, such as school closures, mass vaccination, or the use of face masks. The results of this study suggest that such an approach to pandemic modeling is fundamentally limited. The effects of pandemic control strategies may vary considerably in the presence of other control strategies, or be contingent on the epidemiological characteristics of the particular strain causing the pandemic. There may be a risk of misrepresenting or oversimplifying the effectiveness, and cost-effectiveness, of pandemic control strategies if these strategies are not evaluated in a way that considers the presence of other policies or variable pandemic characteristics.

REFERENCES

- Aleman, D. M., T. Wibisono, and B. Schwartz. 2011. "A non-homogeneous mixing model for predicting pandemic spread". *Interfaces Special Issue on Humanitarian Applications: Doing Good with OR* 41:301–315.
- Beeler, M. F., D. M. Aleman, M. W. Carter, and B. Schwartz. 2011. "Improving Influenza Pandemic Mitigation Policy through Agent-Based Simulation". In *Proceedings of the 2011 Industrial Engineering Research Conference*, edited by T. Doolen and E. V. Aken. Reno, NV: Institute of Industrial Engineers.
- Das, T. K., A. A. Savachkin, and Y. Zhu. 2008. "A large scale simulation model of pandemic influenza outbreaks for development of dynamic mitigation strategies". *IIE Transactions* 40:893–905.
- Durbin, A., A. N. Corallo, T. G. Wibisono, D. M. Aleman, B. Schwartz, and P. C. Coyte. 2011. "A cost effectiveness analysis of the H1N1 vaccine strategy for Ontario, Canada". *Journal of Infectious Diseases and Immunity* 3:40–49.
- Haber, M. J., D. K. Shay, X. M. Davis, R. Patel, X. Jin, E. Wientraub, E. Orenstein, and W. W. Thompson. 2007. "Effectiveness of interventions to reduce contact rates during a simulated influenza pandemic". *Emerging Infectious Disease* 13:581–589.
- Khazeni, N., D. W. Hutton, A. M. Garber, N. Hupert, and D. K. Owens. 2009. "Effectiveness and Cost-Effectiveness of Vaccination Against Pandemic Influenza (H1N1) 2009". *Annals of Internal Medicine* 151:829–840.
- Los Alamos National Laboratory 2006. "EpiSimS Los Angeles Case Study". <http://public.lanl.gov/stroud/LACaseStudy5.pdf> Last accessed March 15, 2011.
- Sander, B., C. T. Bauch, D. Fisman, R. A. Folwer, J. C. Kwong, A. Maetzel, A. McGeer, J. Raboud, D. C. Scales, M. Z. Gojovic, and M. Krahn. 2010. "Is a mass immunization program for pandemic (H1N1) good value for money? Evidence from the Canadian Experience". *Vaccine* 28:6210–6220.
- Valle, S. Y. D., J. M. Hyman, H. W. Hethcote, and S. G. Eubank. 2007. "Mixing patterns between age groups in social networks". *Social Networks* 29:539–554.

AUTHOR BIOGRAPHIES

MICHAEL F. BEELER is a graduate student in Operations Research at the Massachusetts Institute of Technology. He holds a masters degree in Industrial Engineering from the Department of Mechanical and Industrial Engineering at the University of Toronto. He is also a CIHR-STIHR Public Health Policy Fellow, CIHR-PHARE Fellow, and Fulbright Science and Technology Scholar. His research interests are pandemic planning and supply chain management in low-income countries. His email address is michael.beeler@utoronto.ca.

DIONNE M. ALEMAN is an Assistant Professor in the Department of Mechanical and Industrial Engineering at the University of Toronto. Her research interests are medical applications of operations research, specifically, pandemic planning and radiotherapy treatment optimization. She serves as the President of the INFORMS Junior Faculty Interest Group and the Secretary and Treasurer of the INFORMS Section on Public Programs, Services, and Needs, and is a past Chair of the INFORMS Health Applications Section. Prof. Aleman is an Associate Editor for IIE Transactions on Healthcare Systems Engineering and the Journal of Biomedical Data Mining. She is also the Director of the Medical Operations research Lab (morLAB) at the University of Toronto. Professor Aleman's email address is aleman@mie.utoronto.ca.

MICHAEL W. CARTER is a Professor in the Department of Mechanical and Industrial Engineering at the University of Toronto, and Director of the Centre for Research in Health Care Engineering. His research interests are healthcare resource modeling, demand forecasting and scheduling. He is on the editorial board for the Journal of Scheduling and the journal Health Care Management Science, and is an Adjunct Scientist with the Institute for Clinical Evaluative Sciences in Toronto (www.ices.on.ca). Professor Carter's email

Beeler, Aleman, and Carter

address is carter@mie.utoronto.ca.