Published in final edited form as:
J Public Health Manag Pract. 2013 ; 19(0 2): . doi:10.1097/PHH.0b013e3182939a5c.

A Sequential Experimental Design Method to Evaluate a Combination of School Closure and Vaccination Policies to Control an H1N1-Like Pandemic

Kiatikun Louis Luangkesorn, PhD, **Farhad Ghiasabadi, MS**, and **Jagpreet Chhatwal, PhD**
Departments of Industrial Engineering (Drs Luangkesorn and Chhatwal and Mr Ghiasabadi) and Health Policy and Management (Dr Chhatwal), University of Pittsburgh, Pittsburgh, Pennsylvania.

Abstract

Context—During the 2009 H1N1 pandemic, computational agent-based models (ABMs) were extensively used to evaluate interventions to control the spread of emerging pathogens. However, evaluating different possible combinations of interventions using ABMs can be computationally very expensive and time-consuming. Therefore, most policy studies have examined the impact of a single policy decision.

Objective—To apply a sequential experimental design method with an ABM to analyze policy alternatives composed of a combination of school closure and vaccination policies to provide a set of promising “optimal” combinations of policies to control an H1N1-type epidemic to policy makers.

Methods—We used an open-source agent-based modeling system, FRED (A Framework for Reconstructing Epidemiological Dynamic), to simulate the spread of an H1N1 epidemic in Alleghany County, Pennsylvania, with a census-based synthetic population. We used an approach called best subset selection method to evaluate 72 alternative policies consisting of a combination of options for school closure threshold, closure duration, Advisory Committee on Immunization Practices prioritization, and second-dose vaccination prioritization policies. Using the attack rate as a performance measure, best subset selection enabled us to eliminate inferior alternatives and identify a small group of alternative policies that could be further evaluated on the basis of other criteria.

Results—Our sequential design approach to evaluate a combination of alternative mitigation policies leads to a savings in computational effort by a factor of 2 when examining combinations of school closure and vaccination policies.

Conclusions—Best subset selection demonstrates a substantial reduction in the computational burden of a large-scale ABM in evaluating several alternative policies. Our method also provides policy makers with a set of promising policy combinations for further evaluation based on implementation considerations or other criteria.

Keywords

agent-based model; best subset selection method; H1N1 pandemic; simulation

Correspondence: Kiatikun Louis Luangkesorn, PhD, University of Pittsburgh, 1048 Benedum Hall, 3700 O’Hara St, Pittsburgh, PA 15261 (lol11@pitt.edu). .

The authors declare no conflicts of interest.

Agent-based models (ABMs) can capture complex system dynamics arising out of individual behavior and interactions of heterogeneous agents. They are increasingly used to simulate the spread of infectious agents and evaluate mitigation strategies. In addition, ABMs can be used to explore the costs and benefits of various disease prevention, diagnostics, and treatment interventions for even new agents and aid policy makers in forming an overall response policy under resource-limited settings.[1] However, ABMs that capture different aspects of geography and disease spread to simulate the real-life situation can be complex and computationally intensive.

Agent-based models have been used to help public health decision makers and agencies prepare for potential events such as pandemics or bioterrorism attacks by evaluating potential responses such as vaccination policies[2-5] or school closure policies.[6,7] However, these models typically evaluate a only a small number of candidate strategies instead of considering a full spectrum of policy options.[8] The computational burden discourages modelers from evaluating multiple combinations of mitigation strategies. During the time of a pandemic, delays in the translation of results of modeling experiments to relevant policy makers can have a substantial impact on the population health.

While evaluation of multiple strategies can be computationally expensive, well-designed experiments can minimize the effort needed to provide estimates of the performance of alternative policies. Classical designs such as the Hsu[9] method for Multiple Comparisons with the Best can be used to determine the computational effort needed to evaluate a number of alternatives and identify which ones are close to the best with a specified statistical level of precision. In the case of simulation experiments, there are methods that adjust the number of replications required during the middle of the experiment. These sequential procedures iteratively use screening procedures that steadily reduce the number of alternatives that require further simulation until they obtain the final result.[10]

In this work, we evaluate the use of sequential experimental design methods to reduce the computational burden of evaluating multiple intervention policies in a complex agent-based disease transmission model.[11] Specifically, we apply the best subset selection (BSS) method to the FRED (A Framework for Reconstructing Epidemiological Dynamic) simulation model of the spread of H1N1 epidemic in Alleghany County, Pennsylvania, with a census-based synthetic population. The BSS method selects good policies that are close in performance to the best policy.[11] This is in contrast to optimization methods that attempt to find the optimal policy. The BSS method supports health policy makers' desire to identify a set of alternatives instead of a single "best" policy.[12] This allows policy makers to use the results of models even in cases when the models used do not address all of the issues that need to be considered by the policy makers.

Methods

We simulated the spread of influenza and several mitigation policies in Allegheny County, Pennsylvania. We used FRED, developed at the University of Pittsburgh. FRED is a modular, open-source framework that uses agent-based modeling based on a census-based synthetic population that captures the demographic and geographic distributions of the population as well as detailed household, school, and workplace social networks.

We examined a combination of 4 intervention policies: (1) the school closure threshold, that is, if the attack rate in a given school increases a certain threshold, close the school; (2) the school closure period; (3) if the Advisory Committee on Immunization Practices (ACIP) recommendations on vaccine priorities is used; and (4) the vaccination priority rule to be used in the case where 2 vaccine doses are required. The threshold for closing individual schools was set to 3 possible values: 0.02, 0.05, and 0.10 of an individual school population infected. For the school closure period, 3 options were implemented: 5, 10, and 15 days. This corresponds to 1, 2, and 3 weeks of school closure for each instance. On the basis of information published by the Centers for Disease Control and Prevention,[13] the vaccine was modeled in FRED to have the following characteristics: 1 dose is required for those older than 10 years, and 2 doses for those up to 9 years of age. For those 9 years and younger, the second dose is required 28 days after the first dose. Vaccines become available by day 7 of the epidemic and are available as needed thereafter. Vaccine efficacy by age was set as determined in Zhu et al.[14]

Next we examined 2 aspects of the vaccine policy—the choice to follow the ACIP guidelines (yes or no), and the priority rule for the use of second dose. The ACIP priority policy indicates whether or not the ACIP priority guidelines were followed. This includes giving priority in vaccinating persons aged 0 to 24 years, people deemed at risk for complications for influenza, pregnant women, and people older than 64 years.[15] The vaccine priority rule determines the priority of persons getting a second dose versus people getting their first dose. The 4 levels of vaccine priority being examined are as follows: (1) first come first serve; (2) place people getting the second dose at the front of the queue; (3) mix in people getting the second dose with other priority vaccinations randomly; and (4) place people getting the second dose at the end of the queue. The combination of all of these policy options generated 72 potential policy combinations.

We applied a sequential experimental design method to evaluate the above-defined policy options.[11] This takes as inputs 2 parameters: a desired region and an acceptable region (Figure 1). The desired region should represent the difference in performance measures where the 2 policies should be viewed as having practically the same performance. The acceptable region should represent the uncertainty in the simulation results so that if a policy is close to the desired region boundary, it does not matter whether or not it is included in the final results given to the decision maker. These parameters define the policy maker’s indifference zone, that is, a performance region near the best alternative policy where policy outcomes are considered to be effectively equal. Our method runs a small number of simulation replications for all alternatives and then runs additional replications for all policy alternatives that have not been eliminated because of poor performance. Doing so efficiently uses computational resources on good performing policies.

For our analysis, we used the attack rate at the performance measure. The indifference zone was set to 1.0% and the acceptable region to 0.5% difference in attack rate. We simulated 12 initial replications, n_0, of each policy before the method began screening policies. Each replication simulated 100 days from the introduction of the agent.

Results

Following the initial n_0 replications, the BSS method identified 8 alternatives that were close to the best policy for further evaluation. After 6 more rounds of replications, these 8 alternatives were confirmed as the set of alternatives that lay within the desired or acceptable regions. This required a total of 912 replications across the 72 alternatives. We plotted the mean attack rate and variation as determined through the simulation runs of the 72 policies, using the box and whiskers plot in (Figure 2).

The Table summarizes the set of 8 policies across 4 policy dimensions that were found to be significantly better than the others. The policies in the best subset correspond to the policies with a low school closure threshold (0.02), follow the ACIP recommendations on vaccine priority, and have either 10- or 15-day school closure periods. In comparison, the priority scheme for the second dose for those 9 years and younger did not have a significant effect on the simulated attack rate. In particular, these results suggest that because the difference in the 10- and 15-day school closure period is not of practical significance, a policy maker would choose the 10-day closure period after recognizing the additional expense incurred by a 15-day school closure.

We also compared the computational burden of our method with a traditional design of experiment method. The Hsu[9] method of Multiple Comparisons with the Best determines a replication number for all options based on simultaneous confidence intervals. This method can be made to correspond to BSS by setting the minimum significant difference to be the same as the indifference zone used with BSS. Similar to BSS, the Hsu method does not assume prior knowledge of the best policy; instead, the Hsu method compares each policy in terms of its distance from the best and does not predetermine how many (if any) policies would be considered close to the best.

Using the Hsu method, we determine the number of replications that can be expected to be required to construct the 72 simultaneous 100% (1%-0.05%) confidence intervals for the difference between each configuration and the best (lowest attack rate) configuration. Since we set the indifference zone to 1.0 and the acceptable region to 0.5, we will take as a goal to detect a minimum significant difference in attack rate of 1.0 + 0.5 = 1.5%. Following the Hsu method, we find that approximately 23 replications of each alternative policy would be required, resulting in a total of 1656 replications across 72 policy alternatives.

In comparison, when implementing the BSS method, a total of 912 replications were required across 72 policy alternatives, implying 45% reduction in computational burden to achieve the same statistical guarantee of identifying policies that should be included. This savings comes from the fact that after the initial replications and after every round of simulations, the BSS method makes a decision on the remaining alternatives and determines which alternatives require continued analysis based on the specified desired region, acceptable region, and level of confidence.

It should be noted that the replications identified by this use of the Hsu method are an approximation, based on the estimate of the mean and variance of the performance of each alternative over the initial $n_0 = 12$ replications. Therefore, in principle, there is no statistical guarantee that at the conclusion of the 23 replications that the resulting confidence intervals would allow us to correctly categorize each alternative as close or not close to the best policy.

Discussion and Conclusion

We demonstrate an example of using sequential experimental design methods along with an established ABM to examine a combination of different mitigation policies to respond to an infectious disease outbreak. Our method reduces the computational burden of running large-scale ABMs by rejecting poor performing policy alternatives quickly. The savings of 45% in computational burden could lead to a more rapid response to policy makers or evaluation of interactions of policies across more axes even with limited time resources. In addition, since the method generates a set of results, policy makers can then take the identified alternatives and evaluate them on the basis of criteria not addressed in the model. (eg, considering societal cost or social disruption in addition to attack rate).

However, the savings demonstrated are specific to the characteristics of the agent, the intervention, and the set of the policy options that were to be examined. One limitation to this study was a limited range of values that could be covered on each policy axis. In addition, this study examines only the uncertainty due to the Monte Carlo simulation and not parameter uncertainty. This method can be combined with experimental designs that directly address these issues so that these can be examined across multiple policy axes simultaneously.

Another practical approach to evaluate several multiple policy combinations for making real-time decisions is to run hundreds of simulations a priori, as done by the FRED Navigator.[16] However, this approach requires the knowledge of the infectious agent's characteristics, which is not always possible for a novel agent.

Sequential experimental design methods have a promise of reducing the computational burden of running large-scale ABMs. However, to examine the interaction effect of the range of policy options that need to be considered in infectious disease control, savings on the order of 45% such as that observed in this study may not be sufficient. More work needs to be done in exploring other methods for looking at multidimensional policy interactions and how other methods such as stochastic kriging or meta-modeling perform when exploring a large state-space model for infectious disease control.

Acknowledgments

This work was supported by the University of Pittsburgh Models of Infectious Disease Agent Study (MIDAS) National Center of Excellence grant 5U54GM088491-03. The funding agency had no role in the design, conduct, and data analysis of this study.

The authors thank the assistance and advice of John Grefenstette, Shawn Brown, and David Galloway, with the University of Pittsburgh Public Health Dynamics Laboratory, for their explanations of issues involving the use of agent-based models to advise policy makers and their assistance in using the FRED model.

REFERENCES

1. Long, EF.; Brandeau, ML. OR's Next Top Model: Decision Models for Infectious Disease Control. INFORMS; San Diego, CA: 2009. p. 123-138.http://tutorials.pubs.informs.org/ [Accessed November 20, 2009]
2. Lee BY, Brown ST, Korch GW, et al. A computer simulation of vaccine prioritization, allocation, and rationing during the 2009 H1N1 influenza pandemic. Vaccine. 2010; 28(31):4875–4879. [PubMed: 20483192]
3. Lee BY, Brown ST, Cooley PC, et al. A computer simulation of employee vaccination to mitigate an influenza epidemic. Am J Prev Med. 2010; 38(3):247–257. [PubMed: 20042311]
4. Lee BY, Brown ST, Cooley P, et al. Vaccination deep into a pandemic wave potential mechanisms for a "third wave" and the impact of vaccination. Am J Prev Med. 2010; 39(5):e21–e29. [PubMed: 20965375]
5. Lee BY, Brown ST, Bailey RR, et al. The benefits to all of ensuring equal and timely access to influenza vaccines in poor communities. Health Aff (Millwood). 2011; 30(6):1141–1150. [PubMed: 21653968]
6. Lee BY, Brown ST, Cooley P, et al. Simulating school closure strategies to mitigate an influenza epidemic. J Public Health Manag Pract. 2010; 16(3):252–261. [PubMed: 20035236]
7. Brown ST, Tai JHY, Bailey RR, et al. Would school closure for the 2009 H1N1 influenza epidemic have been worth the cost? A computational simulation of Pennsylvania. BMC Public Health. 2011; 11:353. [PubMed: 21599920]
8. Dimitrov, NB.; Meyers, LA. INFORMS TutORials in Operations Research. Vol. Vol 7. INFORMS; Hanover, MD: 2010. Mathematical approaches to infectious disease prediction and control; p. 1-25.http://www.informs.org/Pubs/Tutorials-in-OR/2010-TutORials-in-Operations-Research-ONLINE/Chapter-1 [Accessed November 22, 2010]
9. Hsu, JC. Multiple Comparisons: Theory and Methods. Chapman & Hall; London, England: 1996.
10. Bechhofer, RE.; Santner, TJ.; Goldsman, DM. Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons. Wiley-Interscience; New York, NY: 1995.

NIH-PA Author Manuscript
NIH-PA Author Manuscript

11. Wang, Y.; Luangkesorn, KL.; Shuman, L. Best-subset selection procedure. Proceedings of the 2011 Winter Simulation Conference; Phoenix, AZ. IEEE; 2011. p. 4315-4323.http://www.informs-sim.org/wsc11papers/384.pdf
12. Rubin, P. [Accessed May 29, 2012] K best solutions. OR in an OB World. 2012. http://orinanobworld.blogspot.com/2012/04/k-best-solutions.html
13. Centers for Disease Control and Prevention. [Accessed August 10, 2012] 2009 H1N1 monovalent influenza vaccine dosage, administration, and storage. http://www.cdc.gov/h1n1flu/vaccination/dosage.htm. Published 2009
14. Zhu F-C, Wang H, Fang H-H, et al. A novel influenza a (H1N1) vaccine in various age groups. N Engl J Med. 2009; 361(25):2414–2423. [PubMed: 19846844]
15. Advisory Committee on Immunization Practices. [Accessed August 7, 2012] Prevention and control of influenza with vaccines: recommendations of the Advisory Committee on Immunization Practices (ACIP). 2011. http://www.cdc.gov/mmwr/preview/mmwrhtml/mm6033a3.htm. Published 2011
16. Paparian, J.; Brown, ST.; Burke, DS.; Grefenstette, JJ. FRED Navigator: an interactive system for visualizing results from large-scale epidemic simulations. Paper presented at: 8th IEEE International Conference on eScience; October 8, 2012; Chicago, IL: http://www.psc.edu/index.php/escience-2012-workshop

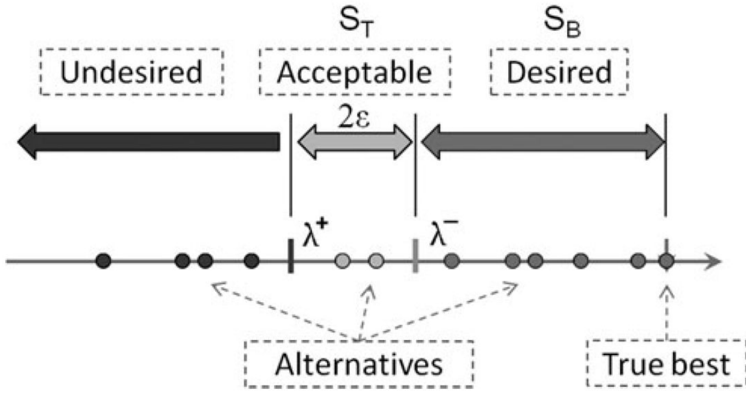

FIGURE 1. Three Regions of Best Subset Selection[a]

[a]The desired region is where policy alternatives must be included in the final result set, policies in the undesired region should not be included, while those in the acceptable region are optional (may or may not be included). The parameters λ^{+} and λ^{-}, which define the boundaries between the regions, should be based on what would be a practically insignificant difference in outcomes, whereas ε, which defines the size of the acceptable region, would be set on the basis of the believed precision of the simulation model.

NIH-PA Author Manuscript

NIH-PA Author Manuscript

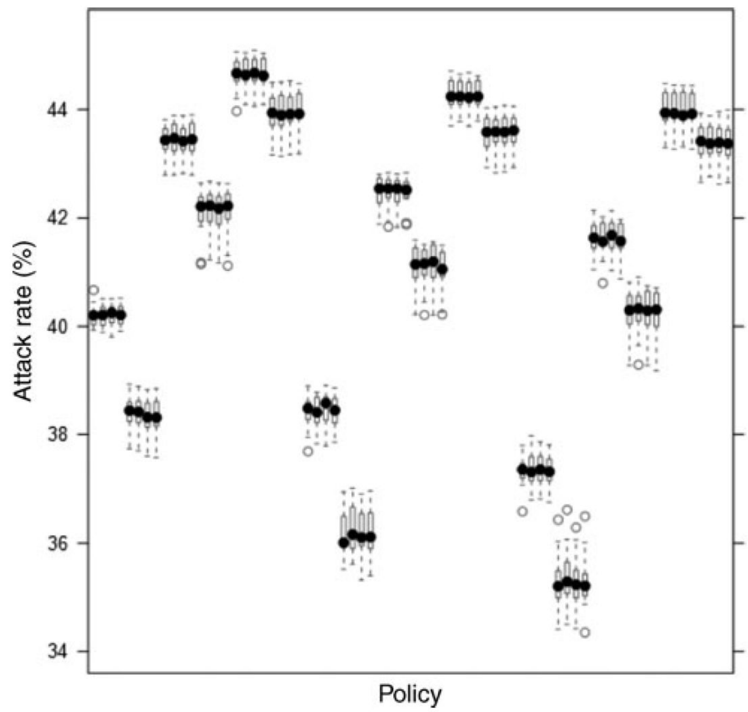

FIGURE 2. Simulated Attack Rate for Each Intervention Policy[a]

[a]For each of the 72 policies, the box and whiskers show the interquartile range and the range of performance for the replications of that system. The dark circles represent the median attack rate for that configuration, whereas open circles represent outliers in reported performance.

TABLE

Policies Identified as a Member of the Best Subset

Closure Threshold	Closure Period, d	ACIP	Second Dose Priority	Attack Rate
0.02	10	Y	0	36.18
0.02	10	Y	1	36.28
0.02	10	Y	2	36.20
0.02	10	Y	3	36.21
0.02	15	Y	0	35.27
0.02	15	Y	1	35.38
0.02	15	Y	2	35.29
0.02	15	Y	3	35.28

Abbreviations: ACIP, Advisory Committee on Immunization Practices; Y, yes.