in silico Public Health:
The Essential Role of Highly Detailed Simulations in Support of Public Health Decision-Making

Bryan Leroy Lewis

Dissertation submitted to the faculty of the Virginia Polytechnic Institute and State University in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Genetics, Bioinformatics, and Computational Biology

Chris L Barrett
Kathleen A Alexander
James R Bohland
Stephen G Eubank

January 19th, 2011
Blacksburg, Virginia

Keywords: public health, simulation, decision support, epidemiology, policy

in silico Public Health: The Essential Role of Highly Detailed Simulations in Support of Public Health Decision-Making

Bryan Leroy Lewis

ABSTRACT

Public Health requires a trans-disciplinary approach to tackle the breadth and depth of the issues it faces. Public health decisions are reached through the compilation of multiple data sources and their thoughtful synthesis. The complexity and importance of these decisions necessitates a variety of approaches, with simulations increasingly being relied upon. This dissertation describes several research efforts that demonstrate the utility of highly detailed simulations in public health decision-making.

Simulations are frequently used to represent dynamic processes and to synthesize data to predict future outcomes, which can be used in cost-benefit and course of action analyses. The threat of pandemic influenza and its subsequent arrival prompted many simulation-based studies. This dissertation details several such studies conducted at the federal policy level. Their use for planning and the rapid response to the unfolding crisis demonstrates the integration of highly detailed simulations into the public health decision-making process.

Most analytic methods developed by public health practitioners rely on historical data sources, but are intended to be broadly applicable. Oftentimes this data is limited or incomplete. This dissertation describes the use of highly detailed simulations to evaluate the performance of outbreak detection algorithms. By creating methods that generate realistic and configurable synthetic data, the reliance on these historical samples can be reduced, thus facilitating the development and improvement of methods for public health practice.

The process of decision-making itself can significantly influence the decisions reached. Many fields use simulations to train and evaluate, however, public health has yet to fully adopt these approaches. This dissertation details the construction of highly detailed synthetic data that was used to build an interactive environment designed to evaluate the decision-making processes for pertussis control. The realistic data sets provide sufficient face validity to experienced public health practitioners, creating a natural and effective medium for training and evaluation purposes.

Advances in high-performance computing, information sciences, computer science, and epidemiology are enabling increasing innovation in the application of simulations. This dissertation illustrates several applications of simulations to relevant public health practices and strongly argues that highly detailed simulations have an essential role to play in Public Health decision-making.

Dedication

I dedicate this dissertation to my wife whose love and support make all I do possible and worthwhile, to my yet-to-be-born daughter whose attendance at my defense was an inspiration, to my brothers, parents, and grandparents for their enduring support and love, and to Daniel, Earnest, Babou, and Galibang Baktar whose openness and generosity in the face of the harsh conditions of their life inspired me to pursue public health and social justice.

Additionally, I wholeheartedly thank and dedicate this dissertation to those who guided my development in the early phase of my academic career. Travis Porco for encouragement and one-on-one mentoring during my master's thesis, for no gain other than furthering the cause of modeling for public health's sake. Reuben Granich for encouragement and the knowledge that enthusiasm and hard work can immeasurably improve the quality and joy of one's work. Strong work! Jennifer Flood for teaching me how to passionately pursue public health in the midst of cold statistical calculations. Joe Eisenberg for introducing me to the world of epidemiological modeling and academic research.

Finally, I am so very grateful for the many personal relationships that sustained me through my years in Blacksburg and dedicate this work to them. Tim Driscoll for companionship during our many sojourns into the wild for relaxation and inspiration. Chris Lasher for his arguable superiority. Andrew Warren for his earnest and steadfast friendship and gift for bringing a smile to anyone's face. My many gracious hosts (Tim and Charley, Andrew and Rachel, Chris, and Andrea) during my frequent trips to Blacksburg, always being welcome makes being on the road more like home. The Dude for his commitment to lackadaisically abiding to his duties of takin' 'er easy for all us sinners. Jenna Swann for her friendship and support of Sarah and I starting on the first day we spent in Blacksburg, VA. Pandapas Wilderness Area for providing joy and stress relief during my many mountain-bike rides down the Beauty trail. Last but not least, the Reverend Steve Sherman who taught me the importance of gratitude and optimism for leading a joyful and long life.

Acknowledgments

I acknowledge and appreciate the support from the many teams whose collaboration enabled the research presented in this dissertation.

Manuscript I
This manuscript was first authored by Betz Halloran and reflects research conducted by teams at Imperial College (Neil Ferguson, Derek Cummings, Christophe Fraser, and Donald Burke), University of Washington (Ira Longini, Shufu Xu, Timothy Germann, Kai Kadau, and Catherine Macken), and Virginia Tech (Stephen Eubank, Bryan Lewis, Anil Vullikanti, Richard Beckman, and Chris Barrett). These efforts were coordinated by the Research Triangle Institute (Diane Wagener and Phil Cooley).

Manuscript II
This manuscript and the research it presents were conducted in collaboration with Ken Kleinman and Allyson Abrams of the Harvard Medical School's Department of Population Medicine who provided expertise on the outbreak detection algorithm and experimental design. Stephen Eubank provided guidance on modeling and experimental design. Additionally Keith Bissett contributed crucial simulation software development.

Manuscript III
This manuscript and the research it presents was conducted in collaboration with teams from the University of Utah (Matthew Samore, Yarden Linvat, and Warren Pettey), the University of Michigan (James Koopman and Yong Yang), and Virginia Tech (Bryan Lewis, Stephen Eubank, Chris Barrett, Madhav Marathe and Keith Bissett). Computing resources were also supplied by the Department of Defense.

I am indebted to my committee members for their patience, helpful advice, and mentorship throughout the development of this dissertation.

Chris Barrett for providing the many fascinating research opportunities, having the vision to create a rich learning environment, and encouraging the exploration of new ideas.

Stephen Eubank for thoughtful comments, fascinating discussions, and his ability to combine complicated research topics and peaceful sail boating.

Kathy Alexander for her mentorship, encouragement, and assistance in setting limits and constraints on a process that can get out of control.

James Bohland for participating on the committee despite a remote location and providing an experienced voice and thoughtful consideration to my research.

Additionally, I acknowledge the support and technical expertise of my colleagues in the Network Dynamics and Simulation Science Laboratory. Madhav Marathe for his advice, interest in my well being as a student, and his ability to coordinate the many people and talents in our group. Keith Bissett for tireless programming and development of cutting edge simulation tools that facilitated all the research presented in this dissertation, yet still having the time and patience to describe even the mundane problems that come with computational research. Paula Stretz for her skill and unending patience setting up the environments and data required to conduct these complex studies. Richard Beckman for his guidance, wisdom, and insights into figuring out the structure of complex data. “Did you plot it yet?” Jiangzhuo Chen for his efficient and vigorous approach to computational problems that enabled so much of this research.

Table of Contents

List of Figures

List of Tables

in silico Public Health: The Essential Role of Highly Detailed Simulations in Support of Public Health Decision-Making

INTRODUCTION

Highly detailed simulations have an essential role to play in supporting public health decision-making. The practice of public health relies on the interpretation of data. Traditionally this data has been deliberately gathered and analyzed to answer specific questions. Alternatively, clever analytic methods have been used to extract as much meaning as possible from existing data, despite its inconvenient structure. While synthetic data generated through simulation cannot replace the knowledge that can be gained from data gathered in the real world, it can enable important public health research in ways that real world data cannot. This dissertation discusses ways that high-performance computing enabled agent-based simulations can be used to create large quantities of highly configurable and realistic data and its essential role in addressing several areas of current public health research.

Well-controlled studies that directly address relevant policy decisions are rarely available, thus necessitating the synthesis of multiple disparate sources of information. Epidemiologic modeling of public health scenarios can support the complexity of combining the myriad associations and point estimates germane to pressing public health problems. This approach has been increasingly relied upon over the past years. This complexity, however, often requires the close collaboration of a few domain experts from the same institution who gather the diverse data, engineer the associations into computer algorithms, execute the computations, analyze the results, and interpret them for decision makers. This approach generates very useful information and has been instrumental to many important policy decisions. However, the next generation of public health research will be aided by methods that focus on the transparent construction of richly detailed layers of synthetic data that can be more openly shared across institutions and information systems. A more data-centric and open environment will readily support interactions between the policy makers with diverse roles, domain experts from a variety of fields, and the synthetic data itself.

EPIDEMIOLOGICAL MODELING

The field of Public Health has experienced significant change in the past decade. Health disparities have been brought into sharper focus through the works of charitable foundations like Medecins san Frontieres, Gates Foundation, Oxfam, UNICEF, etc. The benefits of prevention have gained greater acceptance as the debate over health care reform in the United States has highlighted. Technological advances have increased the ability of the medical profession to combat the pressing health problems of the world through development of new vaccines, treatments, and diagnostic tests. The planned widespread adoption of electronic medical records holds the promise of increasingly

vast arrays of medical information to be analyzed to improve the effectiveness of medical practice. Increased funding to the NIH underscores this increased priority by providing the necessary economic resources to continue advancing biomedical research.

Epidemiologic modeling has played a considerable role throughout the history of public health. From Daniel Bernoulli in the 1760s estimating the impact of smallpox eradication(1) to sophisticated agent-based computer models of pandemic influenza exploring mitigation strategies(2), patterns of disease have been quantitated and numerically analyzed in the quest to ease humankind's suffering. Bernoulli explains, "I simply wish that, in a matter which so closely concerns the well-being of mankind, no decision shall be made without all the knowledge which a little analysis and calculation can provide." (3, 4) This concise description of the mission of public health supports the essential role analytic methods play in public health decision support and explains their role in the evolution of public health over the years.

Most applications of epidemiologic modeling take the form of predicting future outcomes or facilitating complicated calculations. Early mathematical model formulations of population level disease dynamics were described by Kermack and McKendrick in the late 1920's to explain the behavior of epidemics(5). From this foundation their compartmental approach to epidemiologic modeling has been extended and modified to describe just about every disease known to man (HIV-AIDS, TB, influenza, SARS, Ebola, etc). Modeling studies have long been used to advocate for policy decisions or examine particular aspects of disease. Recently they have been directly applied to emerging public challenges to shape policies. In 2001, the foot-and-mouth outbreak in the United Kingdom was the first time that real-time predictive modeling was used to influence crucial decisions and drive the public policy(6). The outbreak of Severe Acute Respiratory Syndrome in 2003 also provided an opportunity for epidemiologic modeling to offer estimates on disease characteristics(7) that strongly influenced policies surrounding a human disease. In the following years, the influenza virus and its potential to cause another serious pandemic as was experienced in 1918, resulted in a considerable level of modeling effort to help make plans to mitigate the anticipated pandemic(2).

With this history of use, epidemiologic modeling has become increasingly accepted among public health policy makers. Indeed, the modeling community has been increasingly turned to and asked to work hand in hand with decision-makers. The ability to build increasingly sophisticated models expands with the continued advances in computer hardware and software. Additionally, there has been a proliferation of high-resolution data gathered from the real world as well methods to access and manipulate these data sources (Navteq, Dun and Bradstreet, Google Earth, OpenStreetMap, etc.). Unsurprisingly, the overabundance of data generated by the various fields of science and business has ushered in an era of data-centric computing, spawning its own discipline in the field of computer science. In this increasingly data-rich environment there is greater need for developing methods to combine and utilize this data in more meaningful ways. Highly detailed agent-based simulations provide a

compelling and satisfyingly transparent method to fuse and create meaning from these disparate sources of data.

ADVANTAGES

Synthetic data is highly adaptable, making it a particularly well-suited medium to bring together different areas of public health research. Public health data is often constrained as the resources available to gather it is often limited. This generally yields information that is specific to a very particular population for several years. Generating this data is expensive and it is often structured to answer a particular question. The adaptability provided by synthetic data can bolster the breadth of more limited data sets or can act as a bridge between one or more different data sets. Additionally the level of detail possible in synthetic data provides points of access for alternative methods. This can facilitate collaborations between the various domain experts.

Synthetic data can provide a canvas for creative and novel experiments. The configurability and relative low cost of creating synthetic data make it a natural foundation for hypothesis generation. Most new tools or methods are initially formed through a process of problem solving, however, many endeavors can be limited a dearth of data. Spiral development methods currently in wide us today require several iterations of testing and further development before a “final” product is ready for release. Synthetic data has the capability of greatly improving this process by enhancing the evaluation process for test users and offering a more diverse set of test cases for the developers. The ability of synthetic data to be highly standardized yet maintain its structure and provide realistic levels of randomness facilitates its integration into this process. The configurability and large volumes of data that are possible can provide a valuable resource for instruction, training, and public health practice evaluation.

LIMITATIONS

While there are many roles for synthetic data there are still significant limitations to its widespread use and adoption. While inviting, this approach is not the most efficient choice for many problems. When appropriate for the problem, applying these methods properly can be computationally intensive and time-consuming. These methods are also specific to the data being fused and require a deep understanding of the limitations of the various data sources. The main difficulty arises from determining the implications of the various constraints of the data sources, however, the flexibility of synthetic data fusion can help make these implications more explicit and provides a framework for a more straightforward integration of the individual data layers. Nonetheless, standards and established methodologies for doing so are still being developed. Increased bandwidth and the emergence of cloud computing(8) can help ease some of the computational constraints, however the complexity inherent with these approaches will remain.

OVERVIEW

The following dissertation will provide examples of three ways that high performance computing enabled agent-based simulations can be used to advance public health research. These examples are drawn from several research projects and are supported by a variety of experiences gained as an epidemiological modeler collaborating with public health decision makers, epidemiologists, statisticians, physicists, mathematicians, computer scientists, software engineers, and other epidemiologic modelers. The first chapter presents work done to directly support public health decision-making with carefully constructed predictive simulation studies. It also details the gradual development of techniques that allows for the integration of models into rapid time-sensitive public health decision-making processes. The second chapter explores the use of the data generated by models to evaluate surveillance systems and outbreak detection algorithms. The third chapter focuses on developing a framework for assessing the ways public health decision-makers use and access data. The varied technologies and methods employed in this research richly demonstrate the benefits of a team science approach.

CHAPTER I. *in silico* Prediction

Informing Public Health Decision-Making with Results from Highly Detailed Simulations

Introduction

Public health decision-making hinges on balancing the costs of each action against its anticipated benefits. The technology, workforce, or know-how might be available to efficiently solve the problem, yet the cost of fully implementing the strategy might be prohibitive. Furthermore, estimating the exact benefits is not trivial since the challenges, the potential solutions, and their real world contexts are so complex. This fundamental challenge has led public health practitioners to develop many methods to assist in these difficult cost-benefit analyses.

Simulations as a Tool for Prediction

A common use of a simulation in this context is to predict and quantify future negative events and evaluate how well different courses of action will mitigate these negative events. For example, the severity of a potential epidemic of infectious disease could be predicted and the impact of various vaccination strategies evaluated for their effectiveness. These studies help decide how to most efficiently distribute resources.

Simulations have been employed to provide predictions of future events to estimate these potential costs for many years, and their capabilities have expanded with advances in the computational fields. Systems of ordinary differential equations are often able to provide estimates of sufficient resolution for decision-makers (9-12). As

the field of epidemiological modeling has advanced, the capabilities of the simulations have improved, as has the expectations of the decision-makers. The level of resolution for the estimates and the complexity of both the disease process and the interventions being evaluated have increased greatly. Different simulation approaches have been adopted to address and support these changes. Agent-based models are often used in place of systems of differential equations. Agent-based models that represent each member of the population individually offer some advantages over less detailed populations that assume higher levels of homogeneity.

Benefits of Agent-based Simulations

The process of overlaying different data sources on an individuated population is more straight forward than a series of one-off calculations to generate the multivariate joint distributions that may be required for a particular analysis. Furthermore, the time-intensive distributions finally generated may not be useful for another set of analyses thus requiring additional effort for modification. Standardizing to a synthetic population of individuals is efficient since most methods for analyzing public health data were initially devised to study data in this form. Additionally, in this form new pieces of information about the disease or population can be readily incorporated.

High-resolution data conveyed in an synthetic individual population ensures compatibility with a wide-variety of epidemilogic modeling approaches. These different approaches may be optimized to operate on data in a particular structure, which is likely to be an aggregation from an individual population. A modular process of data augmentation, aggregation, and simulation, is efficient for a time-sensitive environment.

Data in this form not only expedites the analysis of the the data but can provide a framework for assessing the added value of future data gathering. Often as a public health problem is unfolding the lack of information can arise just as much from not knowing what pieces of information are worth gathering as the inability to gather them. A method to estimate the relative advantages of gathering novel information can play a very important role in answering critical questions in the early stages of an epidemic.

Highly detailed representations of the population not only facilitate the rapid adaption to relevant questions, but further expand the universe of interventions and scenarios that can be represented. Complex interventions and behaviors can be supported by storing characteristics important to the implementation of interventions within each agent. Advances in distributed computing can also support more complex models of disease transmission and progression, allowing for more nuanced representations of disease (asymptomatic infections, a range of severity of disease based on individual characteristics, etc.)(13, 14). Additionally, abstraction of the disease process can allow support for behavioral changes that can dynamically change each agent's exposure to infection. The ability to represent complex disease processes as well behaviors is becoming increasingly necessary to fully support the complexity of issues facing public health decision-makers.

Validation of Simulations

Generally confidence in something new is conferred by its ability to behave in a "trusted way". Frequently for epidemiologic simulations, the trusted behavior is recreating an observation of the real world. Setting aside that this conflates both the underlying process in the real world and the process of observing it (disentangling these two processes can be difficult as addressed in Chapter II), this is generally achievable by reducing complexity through simplifying assumptions. Since even simple simulations with a few variables and no bearing on the system being simulated can be well tuned to recreate any number of real world observations, this simple criterion is not sufficient.

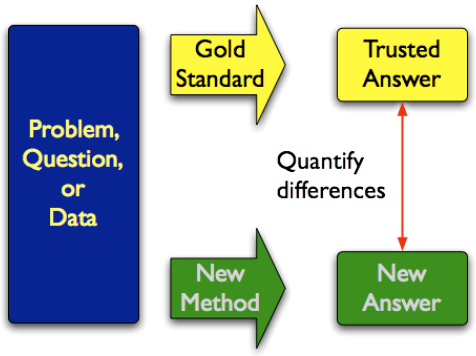

Figure 1. Diagram of process used to develop confidence in a new method

If a simulation is to be relied on for informing important decisions it is natural for the decision-maker to want to have confidence in the model. The process of building confidence is often expressed as the "validity" of the model. There are many criteria for judging model validity in general (15), as well as for complex agent-based models (16). With these concerns taken into account, validity generally refers to the extent that the model or simulation is constructed in a defensible manner and corresponds to the real world. As an imperfect representation of the real world, all simulations cannot be expected to directly mimic the real world system they represent, however, there are several ways in which confidence can be gained.

The manuscript presented in this chapter represents a straightforward approach for developing confidence in the modeling results that is rarely implemented. Three individual teams built models, constructed their interpretations of the scenarios, and simulated results. The independence of the methods demonstrated to the decision-makers that the findings are robust to the specifics of any one method. Since these three results all supported the same conclusion, and in many instances estimated the same relative impacts for the different combinations of interventions, a high degree of confidence was placed in the findings.

Another approach to building confidence in a simulation is to demonstrate multiple levels of correspondence to the real world. If a simulation is constructed to address a particular problem and it is calibrated to a real world system, a correspondence is expected. However, if the construction of the simulation uses high-resolution data and its structure defensibly replicates most of the crucial interactions of

the real world, other aspects of the real world system should be evident in the model. The simulations used for this research demonstrate this level of validity as well.

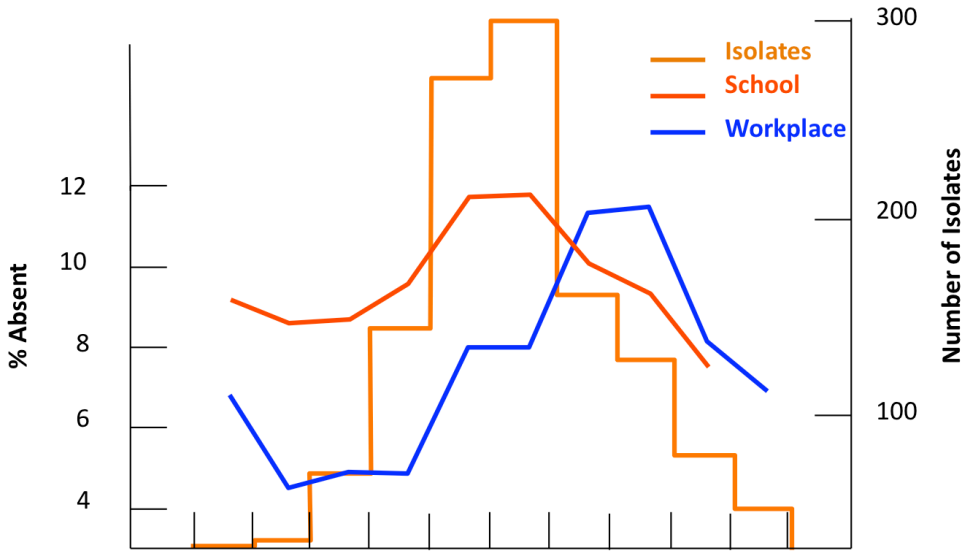

Figure 2. Real world observations from an influenza epidemic in Houston, Texas 1978. School-based absenteeism precedes workplace absenteeism as the epidemic (as measure in influenza positive isolates) peaks.

It has been observed that school-children are on the leading edge of influenza epidemics(17) (Figure 2). Studies in Japan have even shown an improvement in adult populations from vaccination campaigns targeted at school-age children (18). The VBI simulation was constructed to mimic the real world interaction structure. Who interacts with whom was dependent on household structures, age, and income and was not calibrated to match these observations. Analysis of the simulation results (Figure 3a) shows the synthetic school children in the simulation on the leading edge of epidemic. Additionally, when the schools are closed (Figure 3b and 3c), the overall attack rates are greatly diminished and delayed, and it is clear that that the other routes of transmission (college and workplace) are affected as well.

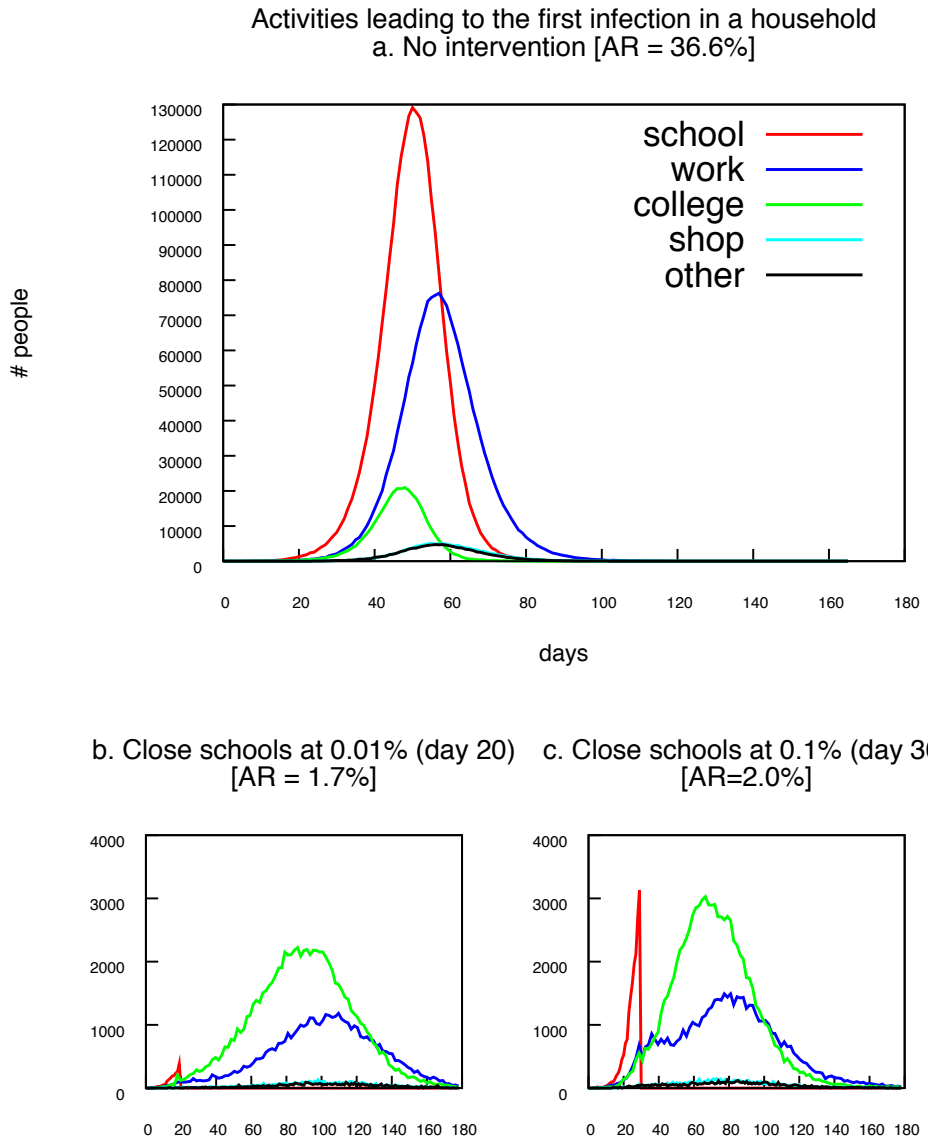

Figure 3. Infection counts of simulated influenza epidemics over time, separated by activity leading to the first infection in a household. A) The unmitigated epidemic infects 36.6% of the population B) Schools are closed 20 days into the epidemic when only 0.01% of the population is infected in a single day C) Schools are closed 30 days into the epidemic when 0.1% of the population is infected on a single day.

Application of Highly Detailed Simulations

Several years ago, as the notion that an influenza virus was likely to cause a pandemic with potentially serious morbidity and mortality gained increasing traction, public health decision makers sought guidance for planning activities. In particular officials at the Department of Health and Human Services (DHHS) and the White House approached the Models of Infectious Disease Agent Study (MIDAS) consortium, of which our group at the Virginia Bioinformatics Institute (VBI), is a member to evaluate the the impact of non-pharmaceutical interventions. Specifically, the decision-makers

were interested in any synergistic effects between the interventions as they were layered on top of eachother.

To support this particular request the MIDAS teams each built independent simulations of influenza spread in a population. All used an agent-based approach but they varied in their actual implementation. The study shows that combinations of simple social distancing can have profound effects on the spread of influenza and should play a role in the mitigation of any influenza epidemic. The study emphasizes the sensitivity of the results to both the timing of the implementation of the interventions as well as the degree of compliance to these interventions. The following manuscript describes the experiences responding to these requests and explains the technical details of the experiments and results.

Additional Studies

As the recent H1N1 pandemic illustrates, despite advances in surveillance, medicine, and public health interventions, infectious diseases can emerge and spread rapidly. The danger posed by diseases of this nature requires careful yet rapid threat assessment and response. The need to refine estimates of the existing situation and what to expect become essential to rational response planning. Epidemiologic modeling has played a role in providing these estimates for the paste several decades. As the sophistication of the models increase, the ability to include them in the real-time decision making loop becomes more challenging because of the increasingly specialized skills required to run, interpret, and analyze the results of these models.

Experience gained during rapid-turnaround research projects led to several refinements to software supporting these projects, and eventually to the development of an environment to support decision-making. This software has been developed and incorporated into real-time decision support both by the developers of the software and by analysts that directly support decision makers. As this software matures it has become clear that the use of this tool can greatly improve the efficiency of existing experts, but can also improve the level of expertise of users that are relatively naive to epidemiologic modeling and analysis.

The adoption of this software and subsequent ease of interpretation of its results has been facilitated by a data-centric approach. Decision-makers are quite familiar with making decisions based on real data. Creating synthetic data with a similar level of resolution and fidelity as data gathered from the real world can allow for direct consumption of modeling results by the decision maker. Additionally, data at this level can be more flexibly fused with a variety of additional data sources, both for incorporation in the models and for further analyses.

A collection of additional studies is detailed in the appendices (B, C, and D) that demonstrate the rapid-turnaround simulation studies also conducted as a part of this dissertation. Appendix B presents results related to an effort to inform the distribution of antiviral medications through the private sector. Appendix C presents results generated to support government analysts in determining the disease characteristics of the newly emergent pandemic H1N1 influenza virus. Appendix D shows the presentation given to

federal public health decision makers to address concerns over a potential resurgence of H1N1 cases following the 2009 holiday season.

Manuscript

The manuscript in Chapter I was published in the Proceedings of the National Academy of Science, in 2008, in volume 105, issue 12, pages 4639-44.

Available online at: http://www.pnas.org/content/105/12/4639.long

CHAPTER II. *in silico* Surveillance

Evaluating Surveillance Systems and Outbreak Detection with Highly Detailed Simulations

Introduction

Public Health decision-making requires accurate and timely situational awareness to be most effective. Determining the actual distribution of a particular disease in a population is challenging. When the timing of interventions is essential, e.g., the mitigation of newly emerging infectious disease outbreaks, this challenge is brought into sharper focus. The existing Public Health infrastructure that gathers this information is often incomplete and can have significant delays. Despite these shortcomings, surveillance systems provide valuable information. Many methodologies and systems have been developed to analyze and interpret this data for routine public health decisions. As a whole these systems act as an effective tool for public health.

The work presented in this chapter integrates disease spread, healthcare seeking behaviors, and a surveillance system into a single simulation. With this capability a framework for evaluating the performance of surveillance systems and the methods based on their data is described. Providing a standardized population on which surveillance system designs can be compared, controls for differences in population characteristics. Generating realistic yet highly configurable synthetic surveillance system data for methods that interpret surveillance data facilitates their evaluation and improvement. This work is the first instance of a large-scale agent-based model being used to evaluate surveillance systems and their dependent methodologies.

Surveillance Systems

Much of the challenge to gaining situational awareness surrounding different public health problems arises from the manner in which the data is gathered. Comprehensive and coherent surveillance networks can be cost-prohibitive, so an ad-hoc surveillance network is relied upon to provide the crucial data needed to track these public health concerns. These systems grew organically from the need to control particular infectious diseases, as a result they are often rigidly focused on a single

disease and the data is often not directly available to all public health officials. The fragmented healthcare system in the United States further complicates the data gathering and integration effort. As a result public health officials devote considerable resources to gathering additional information to enhance situational awareness. Furthermore, additional efforts are needed to properly analyze and interpret this data.

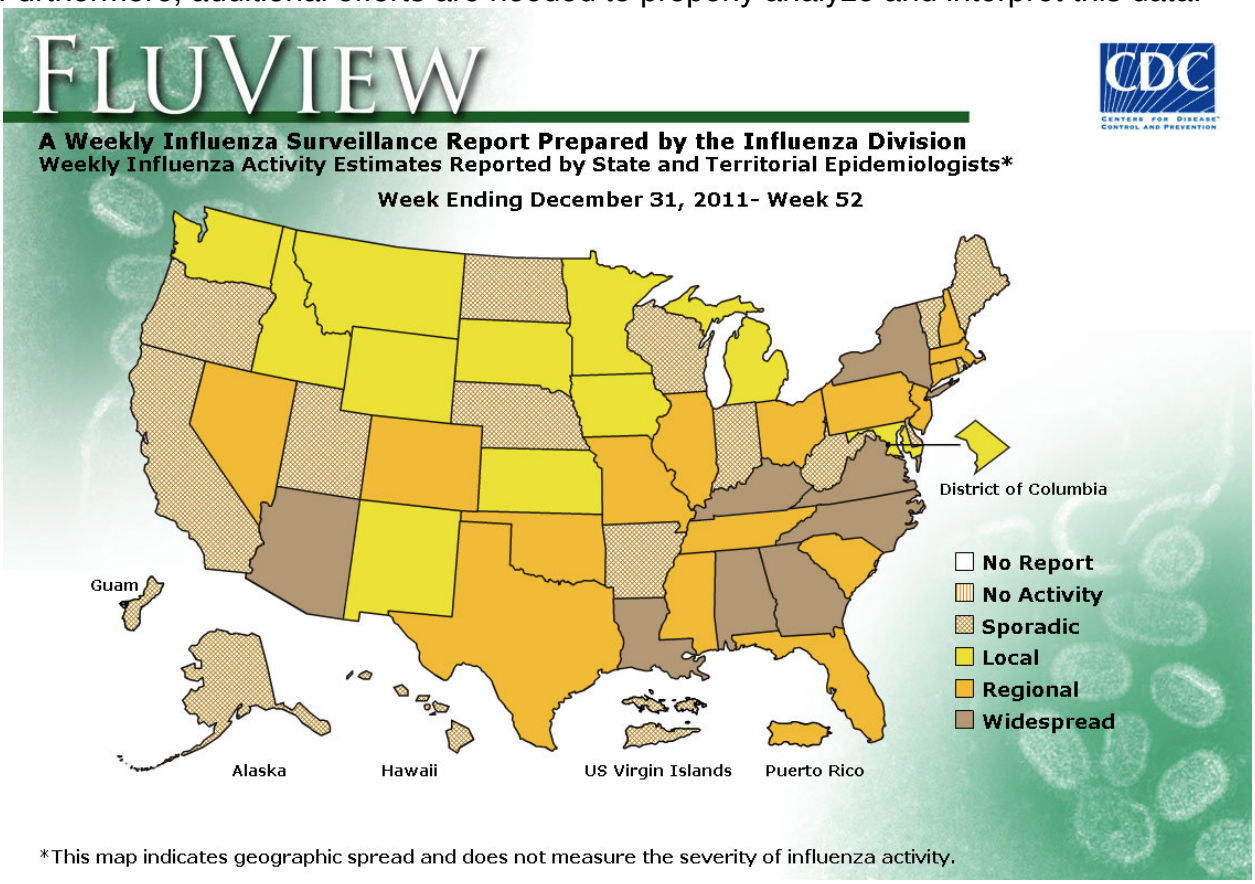

Figure 4. Weekly Influenza Surveillance Report, for December 24-31 (week 52), 2010 (available: http://www.cdc.gov/flu/weekly/usmap.htm; accessed Jan 2nd, 2011)

As an example, consider the influenza surveillance system that operates in the United States. The Centers for Disease Control and Prevention (CDC) and many partners at the global, regional, state, and local level collaborate to gather, share, and analyze a variety of data sources to provide a comprehensive weekly overview (Figure 4) of influenza in the United States and its territories(19). These data come from public health laboratories, healthcare centers, hospital laboratories, death certificates, and state and territorial public health departments. Data of each type is provided by several networks and can be of different resolutions. Even within the networks not all partners are able to provide the exact same resolution (for instance some labs don't have access to technologies to do rapid subtyping). This complex and layered surveillance system carries many different caveats based on the source of the data being interpreted, but nonetheless is a valuable resource for making public health decisions concerning influenza.

The emergence of the pandemic 2009 influenza A (H1N1) virus in 2009 was instructive on the importance of situational awareness. The early identification of this novel pathogen allowed for a series of rapid responses. The virus was isolated and a type specific vaccine was produced and widely distributed to high priority members of the public (health care workers, pregnant women, and the immuno-compromised), greatly reducing the morbidity and mortality of the disease. While vaccine production met with some technical challenges that prevented an optimal distribution of H1N1-specific flu vaccine to the public at large, the H1N1 response was a success. The surveillance system detected anomalous flu cases, tracked the pandemic's growth, and provided useful data for public health officials to base their decisions on (some of these efforts are detailed in Appendices B, C, and D). However, even with the existing infrastructure and deep experience with influenza illnesses, there was still quite a bit of question at the highest levels of public health decision making about the total number of infections caused by the virus in the population(20). Appendix D details a simulation study run to determine the likelihood of an additional wave of infections following the 2009 holiday season. These estimates were found to be very sensitive to the number of previously infected individuals, thus several simulations were needed to estimate what these numbers might be.

The original impetus for this work came from the problem of identifying a potential anthrax outbreak(21, 22). Anthrax, along with many potential bioterror agents, presents a particular challenge since their initial symptoms are indistinguishable from an influenza-like illness (ILI). Setting aside identifying the specific cases of a bioterror event, public health practitioners developed several methodologies designed to identify abnormal patterns in the noisy pattern of normal ILI surveillance data. The goal of these methods was the timely identification of a cluster of cases of concern that could then be investigated.

Detecting Outbreaks

Outbreak detection is one of the core functions of public health. Both the surveillance systems and the methods and activities that interpret the surveillance data are essential to this effort. Outbreak response, as part of routine public health practice, requires timely and focused actions, thus any methods that can alert public health officials to a potential outbreak earlier can have profound public health benefits. In the bioterror context, timeliness is an especially acute concern(23).

Many approaches have been taken to improving the process of outbreak to detection. Many have focused on improving surveillance sensitivity. Gains have been made through direct improvement to laboratory diagnostic tests(24) and clinical case definitions(25, 26). Surveillance data has also been augmented by including syndromic information captured by urgent care centers(27), web-based surveys(28), and by mining Google searches(29) and Twitter updates(30). For bioterror outbreak detection, direct sampling of the environment has been employed for several years(31).

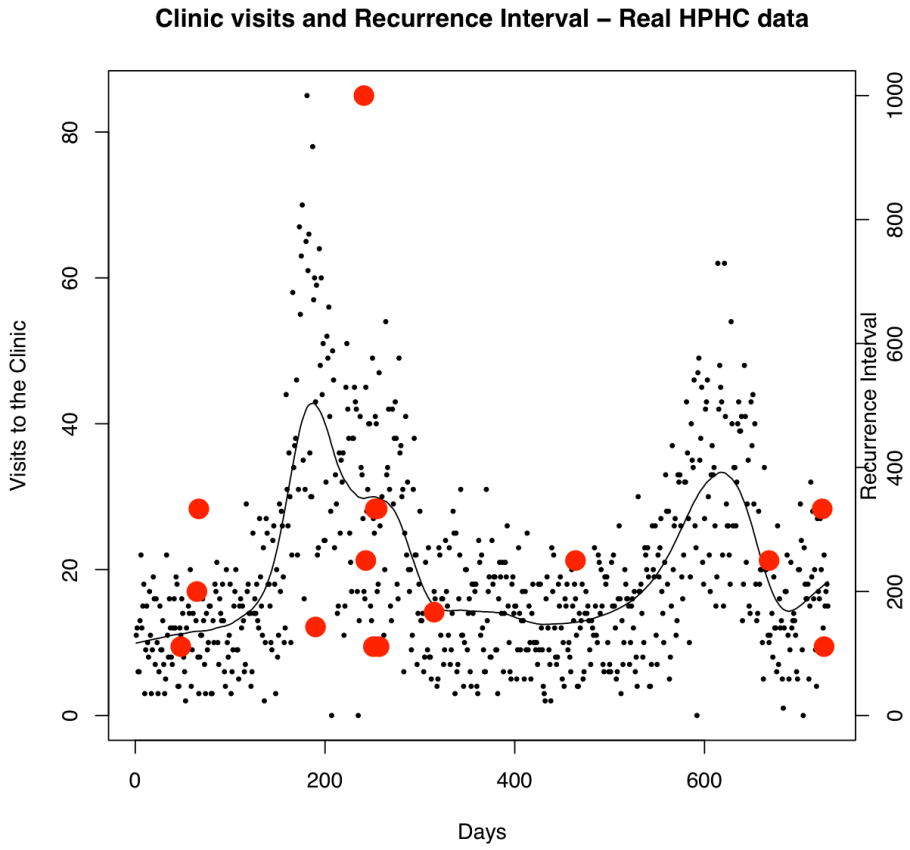

Figure 5. Detection of outbreaks for 2 years of Influenza-like Illness Surveillance data from the Harvard Pilgrim Health Care (HPHC) clinics in the Boston Metro Area. Black dots represent actual counts per day, black line is a smoothed fit of these counts, and the red dots indicate the identification of a cluster and its statistical significance (represented as a recurrence interval, 1/p-value)

Apart from improving the quality or quantity of the data, many sophisticated methods have been developed to analyze these data streams to assist in the identification of anomalous signals (32, 33). Most rely on different ways of aggregating the data being analyzed to generate statistics about what is "normal" and then are able to detect statistically significant anomalies in the signal. For surveillance data that includes location information along with counts, space-time scan statistics have been widely used. A popular implementation of these statistics, SaTScan, was used in this work(22). Figure 5 shows the identification of several clusters by SaTScan on actual ILI surveillance data from the Boston Metropolitan Area.

Evaluation

As with the situational awareness problem itself, the limitations in the data available are the main challenges to evaluation. It is essentially impossible to quantify how many actual outbreaks occurred that should have been detected in any given data

set. Comparing the performance of something as complex as surveillance systems is challenging since the differences between real surveillance systems depends greatly on the population, disease, and the resources available. Controlling for these differences is difficult, especially over the time period needed for proper evaluation.

These limitations on evaluating this crucial public health function hinder improvement of surveillance systems and outbreak detection methods. Since many of the improvements to surveillance must be taken at the local level these limitations are more pronounced since local public health departments have the least resource to devote to evaluation activities. Incomplete evaluation tools also make it difficult for public health officials to allocate their resources most efficiently.

Summary

The manuscript presented in Chapter II addresses these challenges by providing a realistic and configurable virtual world in which to conduct these studies. The methods described directly addresses many of the shortcomings described above. The absolute distribution of disease in the population is known, so there is better knowledge about how many "real" outbreaks occurred. Multiple surveillance systems can be specified simultaneously on the same population and disease process, to provide a controlled environment for unbiased comparisons. Efforts are ongoing to provide these tools as a web-service to extend access to these tools to as wide an audience as possible. By providing tools to assist in the evaluation of surveillance systems and outbreak detection methods it is hoped that overall public health performance can be improved.

Additional Presentations of the Work

The development of this project took several years and was presented in multiple forms and was refined through multiple public presentations. Presentations were conducted at several annual conferences for the International Society for Disease Surveillance (2007-2009) as well as several consortium-wide meetings of the Modeling Infectious Disease Agent Study, which supports this collaborative study. Included in the appendices is a poster presented at the annual American Public Health Association conference in November 2010 in Denver, Colorado (Appendix E). In support of this poster and other public presentations of this work a webpage was developed that provides much of the information presented in this manuscript as well as several movies that illustrate the outbreak detection process in finer detail, the URL and a screen shot are provided in Appendix E.

Manuscript

The manuscript presented in Chapter II was submitted for publication to the American Journal of Public Health on November 30th, 2010.

Available online at: http://db.tt/d5WmeHW

CHAPTER III. *in silico* Instruction

Simulation-Assisted Training and Evaluation for Public Health Decision-Making

Introduction

Public Health practitioners from senior policy-makers to outreach workers in the field operate in the complex and unpredictable environment of the real world. Determining the optimal decisions under these conditions requires extensive knowledge and experience. Instruction for public health practice prepares with this in mind, by equipping the practitioners with the tools to assess their environment and evaluate their decisions. Even with these tools, years of experience are often needed for public health practitioners to be confident in their decision-making. Public Health instruction could benefit from methods that speed up the acquisition of this experience.

Evaluating the process of decision-making scientifically is fraught with many challenges. The decisions intrinsically depend upon their environments, which can be so variable that direct comparison is difficult. If decision-making is constrained to an artificial environment to control for external influences, some of the key aspects of decision-making can be lost. The importance of improving this process has led to a variety of approaches that address these challenges, yet large gaps remain. Methods, like those described in this chapter, that replicate the unpredictability of the real world while remaining replicable and configurable, can fill these gaps.

Public Health Instruction

Demand for Public Health professionals both domestically and worldwide(34) is increasing and higher education institutions are responding by creating more Public Health programs. These new professionals will be increasingly familiar with computational technologies used for instruction(35). To support the increasing demands for trans-disciplinary approaches, well established public health entities, are integrating these technologies into their core curriculums(36). Embracing these technologies is important, however, the need for students to be engaged with the unpredictability of the real world remains(37). These shifts indicate that methods, like those discussed in this chapter, are likely to be accepted and embraced by the next generation of professionals.

Once in practice, professionals have been adopting computer-facilitated methods to further their work. Interactive environments have been developed to educate patients(38-40) on best practices. These environments have also been found to excel at communicating the complexities of different levels of risk(38). Automated interactive environments have also been shown to be useful for training-of-trainer activities(41). As the techniques for integrating these techniques become more widely used and accepted they are likely to become ingrained in public health practice.

Continued education and training to keep abreast of recent changes in the field as well as to gain new skills is essential to good public health practice. For the

established public health professional, use of these techniques can help maintain a fresh skill set. However, experienced health professionals can be more sensitive to issues surrounding the relevance of training to their day-to-day work. Frameworks that enable a highly detailed and interactive approach to instruction, thus providing a more direct link to day-to-day activities, can provide an ideal platform for training.

Evaluation

Evaluation of public health practice is essential to its continued improvement. A large proportion of the public health literature is devoted to evaluating the impact of public health decisions. These efforts have extended to evaluating the processes that led to these decisions, and follow many of the traditional methods of evaluation. To test the efficacy of new decision-making aids interventional studies comparing decisions with and without the aid are conducted(42-44). Qualitative approaches use structured surveys and interviews with experts to better understand their decision-making processes(45, 46). Formalized approaches to studying complex decision making processes like MCDM have been developed and are routinely employed(47, 48). Simulation-assisted techniques are increasingly being adopted to guide decision-making and have been beneficial (49, 50), however, use of simulations as an integral part of the evaluation process has yet to be fully explored.

Role of highly detailed simulations

The methods and process described in this chapter provide a framework for constructing high-impact training exercises and conducting detailed studies of public health decision-making. The main advantages of this approach stem from the highly detailed synthetic information generated by the agent-based simulations at the core of the method. Synthetically created data has the advantage of being highly automatable, thus it can support study designs that require large volumes of similar, yet unique, sets of data. Additionally, the natural stochastic variability within the real world system and its consequences can be well represented by these methods, offering an opportunity to improve instruction on managing the unpredictable. Studies based on these methods, can provide a cost-effective way to evaluate the public health decision making process and improve public health effectiveness.

Summary

The manuscript presented in Chapter III describes the process used to generate a highly detailed data set of sufficient face validity to enable studies of the public health decision-making process. The project required the construction of a large realistic dataset of statewide pertussis cases. This data was generated under several different public health policies to provide a framework for evaluating the decision-making process of public health practitioners. The data structure provided the foundation for the development of a sophisticated interactive software environment (Epinome presented in Appendix G) that provided an extremely data rich environment as well as the ability to capture the information search strategies and policy decisions of the study participants.

The process used to generate the data, how it was used to develop the interactive study environment, and the study it enables are described.

Manuscript

The manuscript presented in Chapter III is intended for submission in Journal of Epidemiology and Community Health in April 2011.

Available online at: http://db.tt/RwA7WfM

INTRODUCTION SUMMARY

These three chapters describe approaches to applying simulations to public health problems. Chapter I describes a direct approach of simulating a complex scenario and making predictions of probable outcomes to be used as the basis for decisions. Chapters II and III describe methodologies that use simulation to enable studies that would previously have been impossible. There are many tools and approaches to rigorously addressing a wide variety (51-58) of problems, however, very few rely directly on simulations. The following manuscripts and conclusion strongly argue that highly detailed simulations have an essential role in supporting public health decision-making.

Chapter I. *in silico* Prediction: Informing Public Health Decision-Making with Results from Highly Detailed Simulations

Modeling targeted layered containment of an influenza pandemic in the United States

M. Elizabeth Halloran*†‡, Neil M. Ferguson§, Stephen Eubank¶, Ira M. Longini, Jr.*†, Derek A. T. Cummings§, Bryan Lewis¶, Shufu Xu†, Christophe Fraser§, Anil Vullikanti¶, Timothy C. Germann , Diane Wagener, Richard Beckman¶, Kai Kadau , Chris Barrett¶, Catherine A. Macken , Donald S. Burke††, and Philip Cooley****

¶Virginia Bioinformatics Institute, Virginia Polytechnical Institute and State University, Blacksburg, VA 24061; ††Graduate School of Public Health, University of Pittsburgh, Pittsburgh, PA 15261; **Research Triangle Institute, Research Triangle Park, NC 27709; §Department of Infectious Disease Epidemiology, Imperial College, London W21PG, England; Los Alamos National Laboratories, Los Alamos, NM 87545; *Department of Biostatistics, School of Public Health and Community Medicine, University of Washington, Seattle, WA 98195; and †Program in Biostatistics and Biomathematics, Division of Public Health Sciences, Fred Hutchinson Cancer Research Center, Seattle, WA 98109

Edited by Barry R. Bloom, Harvard School of Public Health, Boston, MA, and approved January 15, 2008 (received for review July 23, 2007)

ABSTRACT

Planning a response to an outbreak of a pandemic strain of influenza is a high public health priority. Three research groups using different individual-based, stochastic simulation models have examined the consequences of intervention strategies chosen in consultation with U.S. public health workers. The first goal is to simulate the effectiveness of a set of potentially feasible intervention strategies. Combinations called targeted layered containment (TLC) of influenza antiviral treatment and prophylaxis and non-pharmaceutical interventions of quarantine, isolation, school closure, community social distancing, and workplace social distancing are considered. The second goal is to examine the robustness of the results to model assumptions. The comparisons focus on a pandemic outbreak in a population similar to that of Chicago, with 8.6 million people. The simulations suggest that at the expected transmissibility of a pandemic strain, timely implementation of a combination of targeted household antiviral prophylaxis, and social distancing measures could substantially lower the illness attack rate before a highly efficacious vaccine could become available. Timely initiation of measures and school closure play important roles. Because of the current lack of data on which to base such models, further field research is recommended to learn more about the sources of transmission and the effectiveness of social distancing measures in reducing influenza transmission.

Keywords: influenza antiviral agents mitigation prophylaxis social distancing transmission

INTRODUCTION

The on going epidemic of highly pathogenic H5N1 influenza infection in global avian populations has made influenza pandemic preparedness a top public health priority. The interventions being considered fall into two broad classes: medical interventions and non-pharmaceutical interventions (NPIs). Medical interventions include the use of antiviral agents for case treatment, targeted prophylaxis of their known contacts, and prophylactic vaccination. NPIs include social distancing, infection control, and travel restrictions. Social distancing measures include isolation of diagnosed cases, quarantine of households of diagnosed cases, closing of schools, and reducing contacts at workplaces or in the community more generally. Many NPIs were used in U.S. cities during the 1918 pandemic, and appeared relatively successful in some instances, although retrospective assessment is difficult (1-3).

Fundamental to the dynamics of an epidemic is the basic reproduction number, R_0, and the generation time, T_g, of the pathogen (4). R_0 is the average number of secondary cases produced by each primary case at the start of an epidemic in a previously unaffected population, and T_g is the average time between infection of an index case and infection of the secondary cases they produce. Although the R_0 of a future newly emergent influenza strain is unknown, previous estimates are 1.89 from the pandemic in 1968 in Hong Kong (5), and 1.5–1.7 in 1957 in Great Britain (6). The reproductive number of the first wave of the 1918 pandemic A(H1N1) in the United States was estimated as 2–3 (7) and 1.7–2.0 (6). Based on past experience, one might assume for a newly emergent pandemic influenza that R_0 1.7–2.0 and T_g is as short as 3 days. Hence, although an influenza pandemic may be explosive, it is also potentially containable, because reducing transmission by as much as half might achieve an R0 1.

Epidemic models represent a powerful tool for gaining insight into how the dynamics of an epidemic are affected by interventions (8). Small (9, 10) and large-scale (6, 11, 12) individual-based stochastic simulations have previously examined the potential effectiveness of various interventions. However, different research studies seldom examine the same interventions, so results are difficult to compare. In this article, three groups supported in part by the National Institutes of General Medical Sciences MIDAS network coordinated their efforts to use their own stochastic simulation models to examine the same set of intervention strategies. The intervention scenarios and baseline R0 values examined were selected in consultation with government employees working with the Homeland Security Council and the Department of Health and Human Services in the United States, and thus are particularly relevant for the U.S. pandemic plan. One research group is a collaboration of investigators at the University of Washington and Fred Hutchinson Cancer Research Center in Seattle and the Los Alamos National Laboratories (UW/ LANL) (10, 12). One group is a collaboration of investigators at Imperial College and the University of Pittsburgh (Imperial/ Pitt) (6). The third group is at the Virginia Bioinformatics Institute of the Virginia Polytechnical Institute and State University (VBI) (13, 14).

Intervention	Scenario (Compliance %/Ascertainment %)											
	1 Base case	2 30/60		3 60/60		4 60/80		5 90/60		6 90/80	
		Achieved	Compliance	Achieved	Compliance	Achieved	Compliance	Achieved	Compliance	Achieved	Compliance
Symptomatic cases ascertained		60		60		80		60		80	
In ascertained case household											
Threshold	–	1.0		0.1		0.01		0.1		0.01	
Index case treated	–	100	*	100		100		100		100	
Contacts prophylaxed (TAP)	–	100	*	100	*	100	*	100	*	100	*
Home isolation of cases	†		60		60		60		90		90
Quarantine of contacts	–		30		60		60		90		90
School closure	–	100		100		100		100		100	
Threshold	–	1.0		0.1		0.01		0.1		0.01	
Children kept home‡	–		30		60		60		90		90
Workplace distancing	–	50		50		50		50		50	
Threshold	–	1.0		0.1		0.01		0.1		0.01	
Liberal leave	†	100		100		100		100		100	
Community social distancing	–	50		50		50		50		50	
Threshold	–	1.0		0.1		0.01		0.1		0.01	

All numerical values are percentages.
*UW/LANL model assumes 5% stop taking drug after 1 day.
†In all three models, a proportion of symptomatic people retire to home even without intervention.
‡Compliance is % reduction in contacts or contact probabilities outside home.

Table 1. The combined scenarios of targeted layered containment

Intervention Options

We considered a set of interventions consisting of antiviral treatment and household isolation of identified cases, prophylaxis and quarantine of their household contacts, closure of schools, social distancing in the workplace, and social distancing in the community at large. Because these interventions are combinations of targeted and general interventions, we call them targeted-layered containment (TLC) approaches. We examined different levels of ascertainment of symptomatic influenza cases, compliance with the interventions, and cumulative illness attack rate thresholds for initiating interventions.

Initiating the Interventions. Each baseline scenario has a common threshold for all interventions, which varies across the scenarios from 1% to 0.01% cumulative illness attack rate of symptomatic cases.

Ascertainment of Cases. Ascertainment of cases is key for targeted interventions, especially the use of influenza antivirals, case isolation, and quarantine of contacts. Rapid, specific diagnosis will be important. We assume that only 67% percent of influenza infections are symptomatic. We considered two levels of ascertainment of symptomatic influenza cases, namely, 60% and 80%. We assume no asymptomatic influenza infections are ascertained. These levels of ascertainment and pathogenicity correspond to ascertaining 40% and 54% of influenza infections. Interventions within the households of ascertained cases include the following:

- Treatment of ascertained cases. All ascertained cases are treated with one course of antiviral drug for 5 days beginning one day after the onset of symptoms. In the UW/LANL model, 5% of treated cases stop taking the drug after 1 day.

- Targeted antiviral prophylaxis (TAP) of household contacts. All household contacts receive one course (10 days) of prophylaxis beginning 1 day after the onset of symptoms of the index case. In the UW/LANL model, 5% of individuals who receive prophylaxis stop taking drug after 2 days.
- Home isolation of cases. Ascertained cases are isolated in the home, but not isolated from the people with whom they live, with a compliance rate of 60% or 90%.
- Quarantine of household contacts. Household contacts of ascertained cases are quarantined within the home for 10 days with a compliance rate of 30%, 60%, or 90%.

School Closure. All schools, including primary, middle, and high schools, are closed at a particular threshold community cumulative illness attack rate. Once the schools are closed, children are expected to stay at home with a certain compliance rather than to increase community contacts. Compliance is modeled by the reduction in community contacts achieved—assumed to be 30%, 60%, or 90%. In the UW/LANL model, day care centers and small play groups of preschool children are also closed, and the same compliance rates apply. The other two models do not explicitly model day care centers and small play groups.

Liberal Leave Policy. All symptomatic individuals retire to the home from the workplace one day after becoming ill.

Workplace Social Distancing. At a particular threshold community cumulative illness attack rate, workplace contacts are reduced by a certain percent. In the baseline combination scenarios, the workplace contacts are reduced by 50%. Workplaces are not closed. Social distancing in the workplace might eventually be accomplished by staggering the arrivals of workers at work, encouraging people to work at home, or other measures.

Community Social Distancing. Community social distancing represents policies resulting in fewer public activities, such as closing theaters, reducing visits to restaurants, shops, and other public locations, and banning mass gatherings. After a particular threshold attack rate, contacts within the community are reduced by a certain percent, 50% in the baseline combination scenarios. The three models differ in their implementation of community social distancing [see supporting information (SI) Text].

Although we have taken pains to ensure that the models represent the same situations, as described here and in the SI Text, there are subtle model-dependent differences in implementation.

The baseline scenario without intervention, scenario 1, and the five main TLC scenarios are summarized in Table 1. Scenario 2 is the least stringent intervention considered. In scenario 2, interventions are initiated after 1% of the population has developed symptomatic influenza, 60% of clinical cases are ascertained, compliance with quarantine and children staying home after school closing is 30%, and compliance

with isolation is 60%. Scenarios 3 and 5 initiate interventions at an illness attack rate threshold of 0.1% and 60% of cases are ascertained, and differ primarily in assuming 60% versus 90% compliance with interventions. Scenarios 4 and 6 initiate interventions earlier at a threshold of 0.01% illness attack rate and 80% of cases are ascertained, and differ primarily in assuming 60% versus 90% compliance. Scenario 6 is the most stringent TLC intervention considered.

Scenario% compliance/ ascertainment	Intervention threshold, %	R_0 = 1.9 (2.1)			R_0 = 2.4			R_0 = 3.0		
		Imperial	UW	VBI	Imperial	UW	VBI	Imperial	UW	VBI
1	NA	42.4	46.8	44.7	52.4	52.4	51.1	58.8	58.8	56.5
		(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)
2	1	7.3	2.8	3.9	15.5	4.1	9.7	27.4	8.5	20.1
30/60		(104)	(38)	(59.1)	(239)	(61.2)	(140.4)	(421)	(138.4)	(275.1)
3	0.1	1.1	0.31	1.3	4.1	0.41	3.2	18.1	1.03	9.4
60/60		(17.9)	(4.3)	(29.7)	(69.8)	(6.2)	(70.6)	(300)	(17.1)	(189)
4	0.01	0.22	0.04	0.10	0.66	0.05	1.4	10.1	0.12	3.8
60/80		(4.6)	0.76)	(3.2)	(14.2)	(1.1)	(41.7)	(210)	(2.5)	(105.9)
5	0.1	0.17	0.30	1.2	0.66	0.30	2.7	11.5	0.86	6.4
90/60		(3.2)	(4.1)	(28.8)	(14.2)	(5.8)	(59.4)	(192)	(14.2)	(132.2)
6	0.01	0.20	0.04	0.07	0.54	0.05	1.2	5.4	0.11	2.8
90/80		(4.2)	(0.76)	(2.3)	(9.4)	(1.0)	(35.5)	(112)	(2.1)	(80.2)

The Imperial/Pitt model results are based on an average of 10 realizations, the UW/LANL results on an average of 5 realizations, and the VBI results mostly on one realization.

Table 2. Illness attack rates (%) and (antiviral courses per 1,000) using scenarios described in Table 1 in the Chicago population

Sensitivity Analyses.

We undertook the following sensitivity analyses based on scenario 2:

1. Use scenario 2, but vary the percent of workplace and community social distancing between 0 and 50%.
2. Vary the threshold from 0.0001% to 10% community illness attack rate for all interventions in scenario 2.
3. Vary the school closing threshold from 0.0001% to 10% community illness attack rate separately from the 1% threshold for other interventions in scenario 2.
4. Use scenario 2, but with antivirals used only for treatment of ascertained cases, with no prophylaxis of household contacts.
5. Follow scenario 2, but use only nonpharmaceutical interventions, with no antivirals used at all.

The UW/LANL and the Imperial/Pitt groups used their U.S. population models to undertake national-scale simulations of the full TLC as in scenario 2. UW/LANL also explored two further interventions with fewer layers in the U.S. population. One is partial TLC with antiviral treatment of ascertained cases but no prophylaxis, no school closure, and no liberal work leave. The other is with only 50% community social distancing and 50% reduction in long-distance travel and nothing else.

Transmissibility and Case Fatality Ratio.

One uncertainty of a future pandemic strain is how transmissible it will be. It is generally expected that the R0 in a new pandemic will be 2, and previously published articles have explored interventions in this range of R0. Interventions work better at lower R0 values. Here, the focus was on where interventions would break down, which required studying improbably high R0 values. We were interested in examining interventions at R0 values near 2.0, 2.4, and 3.0. At the lower R0, the UW/LANL and VBI models used a value of 2.1, and the Imperial/ Pitt model used 1.9. This is referred to in the text as 1.9 (2.1). A few scenarios at an R0 of 1.6 that are reported in the text.

Another uncertainty of a future pandemic is the case fatality ratio. The estimated 2 in 100 case fatality ratio in the 1918 pandemic (7, 15) is two orders of magnitude larger than the estimated 2 in 10,000 in the 1957 and 1968 pandemics (16). Because the number of deaths that occur will be a fairly linear function of the number of cases and the case fatality ratio, we present only the illness attack rates, and not the number of deaths.

Some Aspects of the Simulation Models

All three models are stochastic, spatially structured, individual-based discrete time simulations. The social structures of the three models are constructed somewhat differently (see SI Text). The UW/LANL basic model is as described in ref. (12), the VBI model in ref. (14), the Imperial/Pitt model in ref. (6). See also refs. (10, 11, 13).

The Chicago area models include the 8.6 million people between latitudes 41.2°N and 42.5°N, and longitudes 87.2°W and 88.5°W, extending slightly across the Wisconsin and Indiana borders. The large-scale simulations of the United States include 281 million people. Each model represents individuals mixing within households, schools, and workplaces, with mixing in the wider community represented differently (see SI Text and Table 3). Transmission can occur in any of the mixing groups represented in the respective models. All three models were calibrated to have age-specific attack rate patterns between those of the 1957 and 1968 pandemics. A few infections are introduced throughout the epidemic in each model as described in SI Text.

In a stochastic, individual-based model, the chance that any susceptible individual will be infected by a contact with an infected person is random and related to the transmission probability for the situation of the contact. Antiviral prophylaxis is assumed to reduce the probability of becoming infected by a contact by 0.3, and if infected, to reduce the probability of developing illness by 0.60. Antiviral treatment or prophylaxis is assumed to reduce the probability of an infected person transmitting by 0.62 (17). Social distancing can lower the number of effective contacts or the transmission probabilities. Many other aspects of each simulation model occur stochastically; for example, whether a person develops symptoms, or whether a person complies with an intervention strategy. In these large populations with the continual seeding of infectives from outside, there is not much variability in the results. The Imperial/Pitt results are based on an average of ten realizations and those of the UW/LANL model on an average of five realizations. The VBI model is much more

computer-intensive than the other two, so the results are based mostly on one realization. The results of a variability study are in the SI Text, SI Fig. 5, and SI Table 5.

Natural History. The natural history within the human host of a future pandemic strain is unknown. All three models have the infectiousness developing before the onset of symptoms, but more of the infectiousness occurs before symptoms in the Imperial/Pitt model than in the other two. This results in a generation time for the UW/LANL and VBI models of 3.2 days, longer than 2.6 days in the Imperial/Pitt model. All three models assume that asymptomatic people are 50% as infectious per contact as symptomatic cases and that the probability of developing symptoms if infected (pathogenicity) is 67% (18).

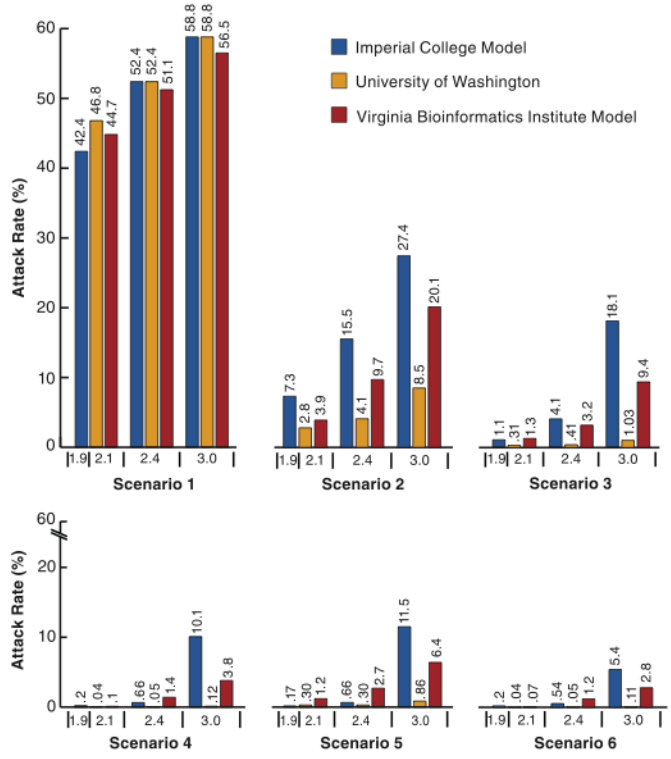

Fig. 1. Influenza illness attack rates for three R_0 values without intervention and with five scenarios of TLC intervention by using the three different models (Chicago population). See Table 1 for a description of scenarios. The R_0 values of 1.9 and 2.1 are considered as a single comparison.

Even in the absence of intervention, all three models assume that clinical disease affects individual behavior. In the UW/LANL and VBI models, 80%, 75%, and 50% of preschool children, school age children, and adults, respectively, with symptomatic influenza withdraw to the home from preschool, school, and work, within the first three days of illness onset. In the Imperial/Pitt model, 90% of symptomatic children do not attend school, 50% of symptomatic adults do not attend work, and community contacts of all symptomatic individuals are reduced by 50%.

Results

Table 2 and Fig. 1 show the results. Increasing attack rates correspond to higher R0 values. In the absence of intervention, the three models produce similar illness attack rates, in the range 42.4 – 46.8% at an R0 of 1.9 (2.1), increasing to the range 56.5–58.8% at an R0 of 3.0. At the lowest R0, in all three models, all five baseline intervention scenarios are effective at reducing the illness attack rates. In scenario 2, at an R0 of 1.9 (2.1), the UW/LANL model achieves a 94% reduction in cases, the VBI model achieves a 91% reduction, and the Imperial/Pitt model achieves an 83% reduction. Although, in scenario 2 at the lower R0, the absolute values of the illness attack rates of the three models range over a factor of 2.6 from 2.8% to 7.3%, the relative effectiveness of the intervention in all three models is high, with the Imperial/Pitt model being the least optimistic. At lower thresholds, higher ascertainment, and higher compliance, the TLC combination is even more effective. At R0 of 1.6 (not in table), the UW/LANL and Imperial/Pitt models produce illness attack rates of 34.7% and 32.0% with no intervention, and 1.9% and 4.5% under the scenario 2 intervention, corresponding to 94% and 85% reductions, respectively.

At the higher R0 of 3.0, the UW/LANL model has an 85% reduction, whereas the VBI and Imperial/Pitt models achieve more modest reductions of 64% and 53%. At the higher R0, the UW/ LANL model is more optimistic than the Imperial/Pitt and VBI models. Part of the difference between the effectiveness of the UW/LANL model and that of the Imperial/Pitt model can be explained by differences in their natural history assumptions. The Imperial/Pitt model has more of the infectiousness earlier so that targeted interventions will have less effect. The difference in the effectiveness of the UW/LANL and VBI models is partly explained by the difference in community social distancing. The VBI model does not close colleges and also has a smaller percentage of the transmissions in the community at large, so that social distancing does not play such a large role. Despite these differences, at the R0 2 or below, the probable range of a pandemic virus, the effectiveness of the interventions in the three models is similar.

	Scenario 1. No intervention			Scenario 2			Scenario 3		
	Imperial	UW	VBI	Imperial	UW	VBI	Imperial	UW	VBI
Illness attack rates	42.4	46.8	44.7	7.3	2.8	3.9	1.1	0.31	1.3
Places									
Home	33.1	39.4	41.1	48.3	58	45.9	50.4	59	36.9
Work	21.8	14.5	28.6	12.9	10	27.8	13.5	10	18.7
School	16.0	18.8	23.3	11.7	11	9.6	9.0	11	2.7
Day care	–	1.1	–	–	0	–	–	0	–
Play group	–	0.8	–	–	0	–	–	0	–
College	–	–	3.3	–	–	12.3	–	–	40.0
Shopping	–	–	2.0	–	–	2.4	–	–	1.0
Neighborhood	–	17.7	–	–	15	–	–	15	–
Neighborhood clusters	–	7.7	–	–	5	–	–	4	–
Other/Community	29.0	0	1.7	26.6	0	2.0	23.8	0	0.8
Totals									
Primary Groups*	70.9	72.7	93.0	72.9	79	83.3	72.9	79	58.3
Community†	29.0	25.4	3.7	26.6	20	4.4	23.8	19	1.8

*Includes home, school, workplace, and for the UW/LANL model, day care and play groups.
†Includes groups subject to community social distancing.

Table 3. Percentage of infections by place and scenario, R0 1.9 (2.1) in the Chicago population

All three models have a large and fairly similar proportion of the infections occurring at home and school (Table 3). The Imperial/Pitt and UW/LANL models have similar amounts of transmissions in the combined school and workplaces, whereas the proportion is substantially higher in the VBI model. The amount in the neighborhoods and neighborhood clusters in the UW/LANL model is similar to that in the community at large in the Imperial/Pitt model. The proportion of infections occurring in the households tends to go up as other sources of infection are closed. Because colleges are not closed in the VBI model, they take on an added importance as a source of infection. The effects of home and school interventions are more robust across the three models than the effects of community social distancing. SI Fig. 6 shows the relative contributions of each activity type in the VBI model to inter-household transmission in the absence of intervention and in scenario 2.

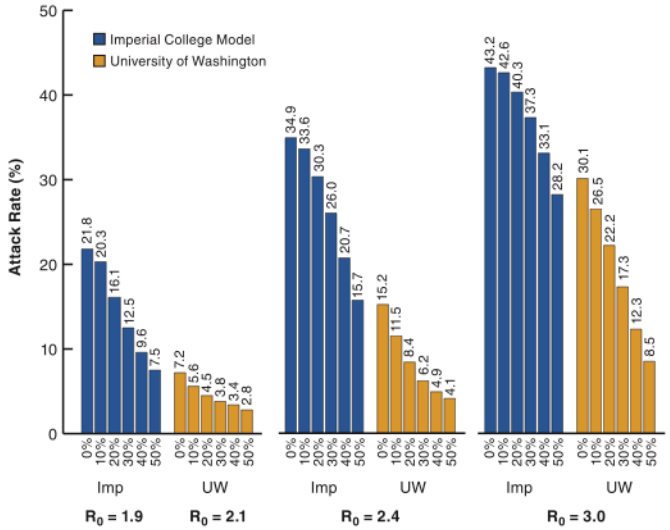

Fig. 2. Sensitivity analysis for workplace and community social distancing. Scenario 2, with community and workplace social distancing being varied between 0% and 50%, and three R_0 values (Chicago population). Only the UW/LANL and Imperial/Pitt models were used. The VBI model is insensitive to changes in this aspect of community social distancing.

Both the UW/LANL and Imperial/Pitt models show an increasing effectiveness in reducing attack rates as community social and workplace distancing increase from 0% to 50%. The corresponding attack rates do not vary in the VBI, so are not shown, but they are the same as in Table 2. The community and workplace social distancing likely play a larger relative role in the Imperial/Pitt model than in the UW/LANL model because the faster natural history makes the targeted interventions based on case ascertainment relatively less effective.

The VBI model is insensitive to the degree of community social distancing, because, as seen in Table 3, only a small proportion of infections occur outside home, school, workplace, and college. It is relatively insensitive to the degree of workplace social distancing, because in that model, workplace social distancing is achieved by reducing the maximum size of workplaces, which does not affect small workplaces.

At R0 of 1.9 (2.1), waiting to implement interventions until reaching a 10% illness attack rate would effect a much smaller reduction in illness attack rates (Fig. 3). However, there is not much improvement by initiating interventions before a threshold

illness attack rate of 0.1%. At R0 of 3, in the UW/LANL model, the lower threshold allows the intervention combination to be highly effective, whereby it achieves only 60% reduction in illness attack rates in the Imperial/Pitt and VBI models. Again, the combination of natural history and community structure of the UW/LANL model make it more optimistic at higher R0 values, whereas all models have similar sensitivity to threshold choice at R0 2.

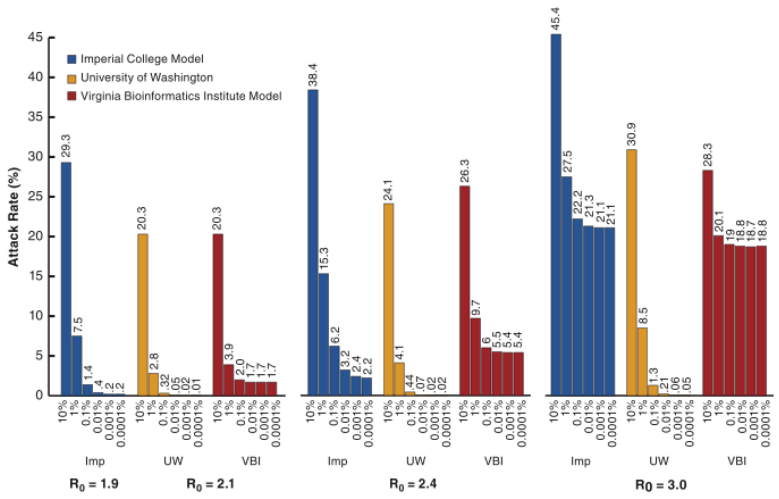

Fig. 3. Sensitivity to changing thresholds for all interventions simultaneously for the three models. Scenario 2 and three R_0 values, with threshold for triggering all measures being varied between 10% and 0.0001% cumulative illness attack rates. Chicago population.

The sensitivity analysis varying only the school closing threshold shows that if schools are closed before the other measures are instituted, the effectiveness of the intervention will be a little greater, but perhaps not enough to warrant the social disruption of early closure of schools (see SI Fig. 7).

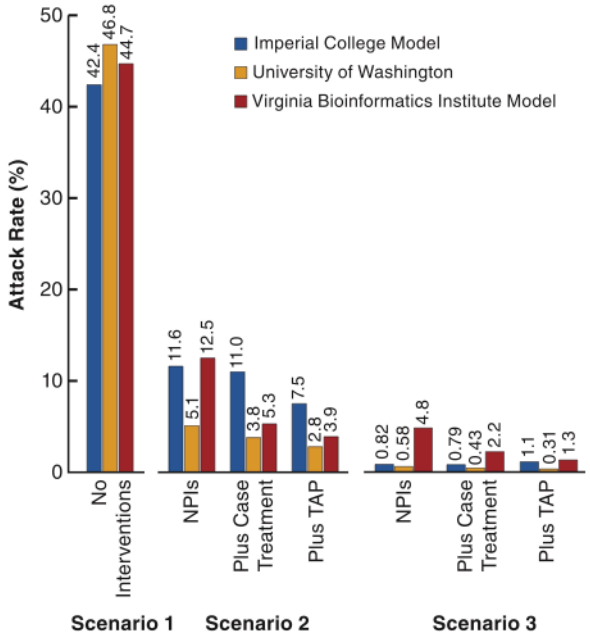

Fig. 4. Comparison of no intervention with intervention scenarios 2 and 3 using just NPIs, NPI with addition of just treatment of ascertained cases (Plus Case Treatment), and NPI with addition of treatment of ascertained cases and targeted antiviral prophylaxis (Plus TAP) of their household contacts. Scenario 1: no intervention; scenario 2: just NPI, with treatment only; with TAP (base case scenario 2); scenario 3: just NPI, with treatment only; with TAP and treatment (base case scenario 3); R_0 of 1.9 (2.1). Chicago population.

In all three models, most of the reduction in the attack rates appears to come from the NPIs (Fig. 4). In scenario 2, the UW/LANL model achieves 94%, the Imperial/Pitt model 88%, and the VBI model 78% of the illness attack reduction with just the NPIs compared with the baseline scenario 2 that uses antiviral treatment and household prophylaxis.

	Illness (infection)	Attack rate, %
Scenario	UW/LANL	Imperial/Pitt
No intervention	47 (70)	42 (63)
Social distancing*	39 (58)	–
Partial scenario 2†	23 (35)	–
Full TLC (scenario 2)‡	0.13 (0.20)	0.30 (0.45)

Threshold is an illness attack rate of 1/1,000 nationally for all interventions except school closure. School closure is implemented locally at the local threshold of 1/1,000 illness attack rate. Otherwise similar to scenario 2 (30/60) when applicable. UW/LANL model $R_0 = 2.1$; Imperial/Pitt model $R_0 = 1.9$.

*Only 50% community social distancing and 50% reduction in long distance travel, nothing else.

†Scenario 2, 50% reduction in long distance travel; but no TAP, treatment only, no school closure, no liberal leave.

‡Scenario 2, school closure at local threshold; 50% reduction in long-distance travel.

Table 4. U.S. national illness (infection) attack rates using three national intervention strategies in the U.S. population models

Table 4 shows results of the UW/LANL and Imperial/Pitt national models of the U.S. population with the full TLC intervention of scenario 2 at the lower R0. In both models, the illness attack rates are substantially reduced. The partial TLC strategy

includes just treatment and isolation of ascertained cases without prophylaxing and quarantining contacts, closing schools, or recommending liberal leave from work for all symptomatic cases. In a third scenario, there is a 50% reduction in community social distancing, such as closing theaters or reduced activities in public places, and a 50% reduction in long-distance travel, not even closing schools. Although the partial TLC strategy still can cut the attack rates in half, the intervention with just community social distancing and 50% reduction in long-distance travel has a much smaller effect, 17% reduction in illness attack rates.

Discussion

Using three different models, we have examined targeted layered containment strategies based on social distancing, rapid case ascertainment, and targeted prophylaxis that, in theory, might be effective in reducing transmission of pandemic influenza. Timely intervention reduces the final number of influenza illnesses.

Especially at values of R0 2 or below, the more probable values for a pandemic strain, the interventions are similarly, although not identically, effective in all three models. At the lower R0, all three models show considerable effectiveness of the suite of NPIs. School closure plays an important role in all three models.

The policy implications have two main aspects. The first is how these results can inform pandemic planning now. If one could achieve these levels of compliance, ascertainment, and social distancing, then there would be a possibility of considerably mitigating a pandemic until a vaccine were available. However, whether the ascertainment and compliance levels modeled here are realistic has yet to be demonstrated. Whether public health officials would actually choose to implement such measures will eventually depend on the lethality and transmissibility (R0) of the pandemic strain. Flexibility in the response plans for different eventualities will be important.

The second aspect for policy is the need for further field research to quantify the natural history of influenza, the sources of influenza transmission, and the feasibility and effectiveness of social distancing measures. Further understanding of the contact structures, such as workplaces and schools, their contribution to the overall transmission of influenza, and how amenable they are to social distancing measures are central to judging which social distancing measures would be effective and worth the social cost.

We caution against overinterpretation of the modeling results, even where the three models suggest similar effectiveness of interventions. Because of the uncertainties in the models, the results need to be viewed more as helping to structure thinking about pandemic planning, rather than being predictive of the precise effectiveness of different policies.

Other simulation results (6, 10, 12) have demonstrated that use of even poorly matched, low-efficacy vaccines would greatly enhance the effectiveness of other intervention measures. Thus, the development and stockpiling of vaccines should be a high priority. When the next pandemic unfolds, it will be important to have the capability

to implement real-time surveillance and epidemiological analysis, including characterizing the new virus, predicting the epidemic trajectory, and if necessary, refining intervention strategies.

References

1. Bootsma MCJ, Ferguson NM. The effect of public health measures on the 1918 influenza pandemic in U.S. cities. Proc Natl Acad Sci USA. 2007;104(18):7588-93.
2. Hatchett RJ, Mecher CE, Lipsitch M. Public health interventions and epidemic intensity during the 1918 influenza pandemic. Proc Natl Acad Sci USA. 2007;104(18):7582-7.
3. Markel H, Stern AM, Navarro JA, Michalsen JR, Monto AS, DiGiovanni C. Nonpharmaceutical influenza mitigation strategies, US communities, 1918-1920 pandemic. Emerging Infect Dis. 2006;12(12):1961-4.
4. Fraser C, Riley S, Anderson RM, Ferguson NM. Factors that make an infectious disease outbreak controllable. Proc Natl Acad Sci USA. 2004;101(16):6146-51.
5. Rvachev RA, Longini IM. A mathematical model for the global spread of influenza. Mathematical biosciences 1985;75(1):3-22.
6. Ferguson NM, Cummings DAT, Fraser C, Cajka JC, Cooley PC, Burke DS. Strategies for mitigating an influenza pandemic. Nature. 2006;442(7101):448-52.
7. Mills CE, Robins JM, Lipsitch M. Transmissibility of 1918 pandemic influenza. Nature. 2004;432(7019):904-6.
8. Influenza CoMCCfP. Modeling community containment for pandemic influenza: A letter report: National Academies Press; 2007.
9. Longini IM, Halloran ME, Nizam A, Yang Y. Containing pandemic influenza with antiviral agents. American journal of epidemiology. 2004;159(7):623-33.
10. Longini IM, Nizam A, Xu S, Ungchusak K, Hanshaoworakul W, Cummings DAT, et al. Containing pandemic influenza at the source. Science. 2005;309(5737):1083-7.
11. Ferguson NM, Cummings DAT, Cauchemez S, Fraser C, Riley S, Meeyai A, et al. Strategies for containing an emerging influenza pandemic in Southeast Asia. Nature. 2005;437(7056):209-14.
12. Germann TC, Kadau K, Longini IM, Macken CA. Mitigation strategies for pandemic influenza in the United States. Proc Natl Acad Sci USA. 2006;103(15):5935-40.
13. Eubank S, Guclu H, Kumar VSA, Marathe MV, Srinivasan A, Toroczkai Z, et al. Modelling disease outbreaks in realistic urban social networks. Nature. 2004;429(6988):180-4.
14. Lewis B, Beckman R, Kumar V, Chen J, Stretz P, Bissett K, et al. Simulated pandemic influenza outbreaks in Chicago Technical Report NDSSL-TR-07-004. 2007.
15. Murray CJL, Lopez AD, Chin B, Feehan D, Hill KH. Estimation of potential global pandemic influenza mortality on the basis of vital registry data from the 1918-20 pandemic: a quantitative analysis. Lancet. 2006;368(9554):2211-8.
16. Kilbourne ED. The Influenza Viruses and Influenza Academic Press. 1975.

17. Yang Y, Longini IM, Halloran ME. Design and evaluation of prophylactic interventions using infectious disease incidence data from close contact groups. J Royal Statistical Soc C. 2006;55(3):317-30.
18. Halloran ME, Hayden FG, Yang Y, Longini IM, Monto AS. Antiviral effects on influenza viral transmission and pathogenicity: observations from household-based trials. American journal of epidemiology. 2007;165(2):212-21.

Author contributions: M.E.H., I.M.L., S.X., and C.A.M. designed research; M.E.H., N.M.F., S.E., D.A.T.C., B.L., S.X., C.F., A.V., T.C.G., R.B., K.K., and C.B. performed research; M.E.H., N.M.F., S.E., I.M.L., D.W., and P.C. analyzed data; and M.E.H., N.M.F., I.M.L., and D.S.B. wrote the paper.
The authors declare no conflict of interest. This article is a PNAS Direct Submission. Freely available online through the PNAS open access option.
‡To whom correspondence should be addressed. E-mail: betz@u.washington.edu.
This article contains supporting information online at www.pnas.org/cgi/content/full/

ACKNOWLEDGMENTS. We thank Richard J. Hatchett and Rajeev V. Venkayya for formulating scenarios of potential interest to the White House Homeland Security Council; Karla Atkins, Keith Bisset, Jiangzhou Chen, Laxminarayana Ganapathi, Achla Marathe, Madhav Marathe, Henning Mortveit, Douglas Roberts, and Paula Stretz (all VBI model) and Simon Cauchemez (Imperial/Pittsburgh model) for helping in developing the original models; and Irene A. Eckstrand for her support and encouragement. This work was supported in part by National Institute of General Medical Sciences MIDAS network Grants U01-GM070749, U01-GM070694, U01-GM070698, and U01-GM070708.

CHAPTER II. *in silico* Surveillance: Evaluating Surveillance Systems and Outbreak Detection with Highly Detailed Simulations

in silico Surveillance: Evaluating Outbreak Detection with Simulation Models

Bryan Lewis, Ken Kleinman, Allyson Abrams, and Stephen Eubank

ABSTRACT

OBJECTIVES Design, implement, and test a flexible methodology for generating detailed synthetic surveillance data that provides realistic geographical and temporal clustering of cases and use to evaluate outbreak detection protocols.

METHODS A detailed representation of the Boston area was constructed, based on data about individuals, locations, and activity patterns. Influenza-like illness (ILI) transmission was simulated, producing 100 years of *in silico* ILI data. Six different surveillance systems were designed and developed using gathered cases from the simulated disease data. Performance was measured by inserting test outbreaks into the surveillance streams and analyzing the likelihood and timeliness of detection.

RESULTS Detection of outbreaks varied from 21% to 95%. Increased coverage did not linearly improve detection probability for all surveillance systems. Relaxing the decision threshold for signaling outbreaks greatly increased false-positives, improved outbreak detection slightly, and led to earlier outbreak detection.

CONCLUSIONS Geographical distribution can be more important than coverage level. Detailed simulations of infectious disease transmission can be configured to represent nearly any conceivable scenario. They are a powerful tool for evaluating the performance of surveillance systems and methods used for outbreak detection.

INTRODUCTION

Detecting outbreaks is a crucial task for public health officials. Distinguishing between truly significant outbreaks and normal disease incidence and to do so as early as possible is a laborious and imperfect science(1). Techniques and methodologies have been developed to screen the increasingly large volumes of data useful to the task (2, 3), yet despite this progress, gaps in performance remain. We present an agent-based simulation methodology to address these challenges.

Developing and evaluating outbreak detection is challenging for many reasons. A central difficulty is that the data used to "train" detection algorithms are unique and relatively brief historical samples and thus do not represent the full range of possible background scenarios. The same dearth of historical data complicates evaluation. In systems where only a count of cases is provided, plausible synthetic data are relatively easy to generate, and can aid in development and evaluation. When evaluating a surveillance system with more precise information, notably detailed geographic locations, simple approaches to generating hypothetical case counts are not plausible. Increasingly realistic simulations of infectious disease spread in highly detailed synthetic populations have emerged in recent years (4-9). Agent-based simulations can represent real-world populations and the day-to-day processes that determine disease-spread and health care seeking behavior. Combining these capabilities with detailed knowledge of a surveillance system allows for the construction of plausible synthetic surveillance data streams. These *in silico* surveillance data streams can be configured to represent nearly any conceivable set of scenarios, making them a powerful tool for studying surveillance systems and outbreak detection algorithms.

We use a novel highly detailed agent-based model, which produces realistic geo-spatial and temporal clustering of influenza-like illness (ILI), to create a library of daily zip code counts from 100 influenza seasons. From this synthetic data, we construct six different surveillance streams that include coverage patterns and health-care-seeking behaviors. For evaluation, we insert artificial outbreaks into the synthetic surveillance data streams at different locations and times and measure the performance of an outbreak detection protocol based on scan statistics. This approach is novel, and represents (to our knowledge) the first time a large-scale agent-based model has been used to evaluate a spatio-temporal outbreak detection tool.

METHODS

Synthetic population creation

The methods used to generate a dynamic and highly detailed synthetic population have been described in detail previously (4, 10). Real-world data and behavioral models are used to generate data sets containing the location of every individual in the population at every moment in the day. Census data and household location data are used to assign individuals to households. The activity profiles, which determine locations throughout the day, were drawn from large surveys and are assigned to each individual based on household and individual demographics, ensuring that individuals perform

activities and in locations appropriate to their age and household structure (11, 12). An iterative fitting process is used to constrain the distances traveled to fit distributions from the activity surveys. This proto-population is then updated with additional information specific to the context of the study being performed, e.g., zip code locations were calculated for all individuals to determine their membership in the various synthetic surveillance systems.

Disease modeling

Influenza and ILI-causing pathogens are transmitted through aerosol contact routes. A discrete event simulation engine simulates the spread of these diseases through the simulated population based on contacts between uninfected and infectious individuals (13, 14). Figure 1 illustrates the structure for the ILI disease model, which is an elaboration on the classic SEIR (Susceptible, Exposed, Infectious, and Recovered) model that accommodates multiple manifestations, asymptomatic infections, and a temporary Recovered state. The detailed disease model (see supplemental materials) used in the simulation includes a number of parameters taken from the literature to represent the broad spectrum of influenza-like illnesses (15-17). Here, the probability of transmission is a function of the duration of contact, symptom severity, type of treatments either individual may have taken, and level of susceptibility of the uninfected individual. For example, if a person, asymptomatically infectious with manifestation 1, was in contact with an individual who was uninfected but had increased susceptibility due to the point in the season, the probability of transmission would be scaled by a level of infectiousness (0.5, due to lack of symptoms), susceptibility (1.15, due to season), and not influenced by treatment (no treatment because no symptoms). Under these conditions and 12 hours of interaction between these individuals the calibrated probability of infection within the simulation would be 0.008322.

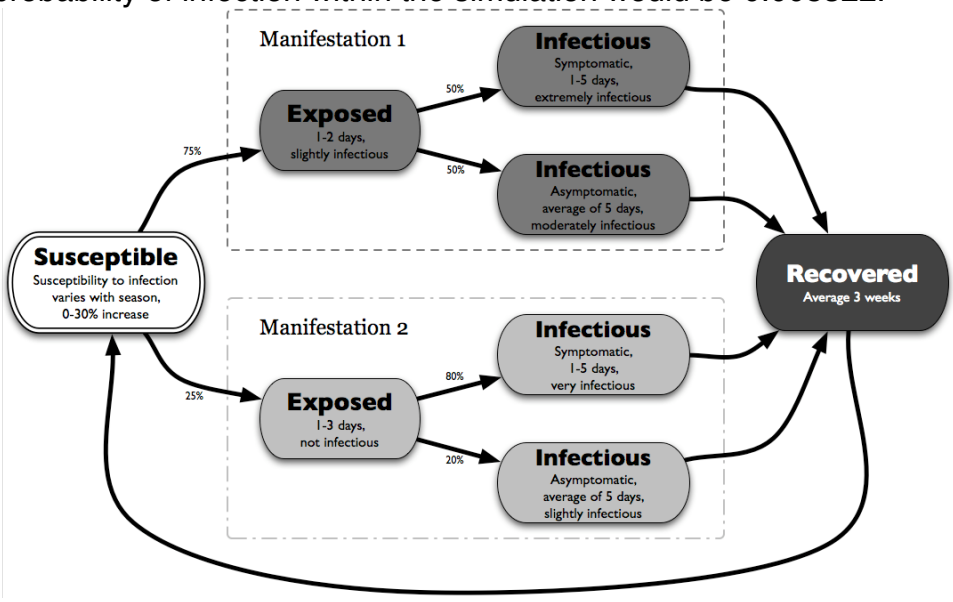

Figure 1. Finite state machine representation allows for a flexible representation of a disease process. Each state determines the duration in that state, level of symptoms, susceptibility, and infectiousness. For example, susceptible individuals have a 75% chance that upon exposure to an infectious contact and successful infection they will transition to the latent Exposed state in ILI Manifestation 1; after 1 to 2 days they will then transition to one of the more Infectious states with a 50% chance of having symptoms.

The characteristic seasonal pattern of ILI infections (Figure 2) was generated. First, an ILI disease was spread in the simulated population, and calibrated to a level that maintained infection at a stable level over several years. Seasonal effects were added by gradually increasing the overall susceptibility to infection of the population in a controlled fashion. The start of this change occurred at a "season onset" date selected at random from a distribution based on historical CDC data (18). Year-to-year variability in season onset and pattern was simulated by choosing from among ten templates inferred from the historical shapes of influenza incidence. Stochastic variability inherent in the system introduced variations between seasons such that even realizations based on the same template differ dramatically.

Care seeking modeling

Health care seeking behavior was modeled by combining the marginal delay from symptom onset to health care seeking distribution (19) and day-of-the-week distribution calculated from Harvard Pilgrim Health Care (HPHC) clinic data for ILI. The resulting joint distribution determines the distribution of the number of days following symptom onset that an individual will wait before seeking health care, given the day of the week. For example individuals who would seek care and start experiencing symptoms on a Monday have a 37% chance of seeking care that Monday, but only an 8% chance of seeking care two days later on Wednesday; whereas if the symptoms start on a Saturday there is a 25% chance they will seek care on that Saturday and a 15% chance they will wait two days till Monday. The probability an affected symptomatic individual seeks health care, is diagnosed correctly, and is reported in the surveillance system can be adjusted to calibrate to a specific real-world surveillance system, or can be parameterized based on data. For example, the effects of a higher co-pay could be introduced to reduce the probability of seeking care, or the probability of correct diagnosis could be reduced to reflect the difficulty of diagnosis.

Surveillance System

Six different surveillance systems were designed. The proportions of Harvard Pilgrim Health Care (HPHC) members living in each zip code in the Boston metro area were determined. Of the 755 zip codes in the synthetic data region, 210 contained the 246,606 members (ranging from 64 to 58,857 people in a single zip code), which is 6% of the entire synthetic population. The selection was done by household and was used as one of the synthetic surveillance systems (Natural 6). The coverage level was doubled (to ~12%) and tripled (to ~18%) while maintaining the same surveillance geography (Natural 12 and 18). The same three coverage levels were combined with a "Uniform" surveillance geography that drew members from the same 210 zip codes, but with an equal proportion of each zip code as members. Thus, the six surveillance

systems evaluated consist of combinations of the two geographic distributions (Natural and Uniform) and the three coverage levels (6%, 12%, and 18%).

Artificial Outbreaks

To test the performance of the different surveillance systems, artificial outbreaks were randomly inserted into the simulated ILI background signals. Each outbreak consisted of cases inserted into the surveillance stream (not subjected to care-seeking behavior modeling) over 14 days, and follows the temporal pattern "2,0,2,0,2,2,0,4,2,0,2,2,0,2" cases per day. This temporal pattern is the same for all insertions. These numbers were scaled with the coverage level of the surveillance system being tested (multiplied by 2 at the 12% level and by 3 at the 18% level). This artificial outbreak was meant to simulate a small concentrated outbreak of influenza-like illness and was based on experience from the Massachusetts Department of Health. The primary location for insertion was randomly chosen from all zip-codes in the population and two additional locations near the primary location were also chosen. The inserted cases were then distributed at random to these three locations.

Outbreak Detection

We assessed the ability to detect the inserted outbreaks using a space-and-time permutation scan statistic (20, 21). Cases from the synthetic population and from the artificial outbreak were assigned a geographic location based on the centroid of their home zip code. The approach finds the cluster least likely to have occurred by chance, where a cluster consists of 1-14 days in any fixed circular region, and assigns a p-value to this cluster using Monte Carlo techniques. This process is repeated daily as the data accrues. The scan implemented uses 90 days of historical data to establish "normal" patterns.

Evaluation of Outbreak Detection

Outbreaks were combined with the simulated ILI background by adding the case counts from the artificial outbreak to the appropriate zip code case count on the appropriate day. The outbreak detection protocol was performed for every day in which inserted cases could have been included in a 14-day putative cluster. The scan statistic was considered to have "signaled" if the p-value fell below a given decision threshold. In this study eight different decision thresholds ($p = 0.033, 0.01, 0.0075, 0.005, 0.004, 0.003, 0.002, 0.001$) were evaluated.

If any scan signal with a p-value at or below a given decision threshold overlapped one of the locations of the inserted outbreak, the inserted outbreak was considered to have been detected at that decision threshold. These are true positives. After attempting to detect the outbreak, the inserted cases were removed before the next outbreak was inserted and its detection attempted. Twelve test outbreaks were inserted in each of the 100 simulated ILI seasons on randomly chosen dates.

This outbreak detection protocol was also performed for the one-year simulated ILI counts without any inserted outbreaks. Any outbreaks signaled in the simulated ILI background data were considered false positives. Any identified signals that overlapped in time and space were "collapsed" into a single "cluster" and designated a false-positive.

Computations

The simulations of ILI disease and health care seeking behavior were performed by EpiSimdemics (14, 22) on an SGI cluster with 96 compute nodes, each with 2 Intel Quad-Core Xeon E5440 processors. The SaTScan analyses were conducted on the same cluster. In total 36,500 scans (100 years * 365 days) with no inserted cases were performed, and 32,400 scans (100 years * 12 insertions * 27 days) were performed for inserted outbreak detection. These simulations and analyses used 84,000 compute hours and required 187 GBs of disk space. Outbreak detection analysis was conducted using SaTScan version 7.0.3 (from www.satscan.org) and R version 2.9.0 (from r-project.org). The analysis package, including code to automate the SaTScan outbreak detection and analysis, will be made available as an R package.

RESULTS

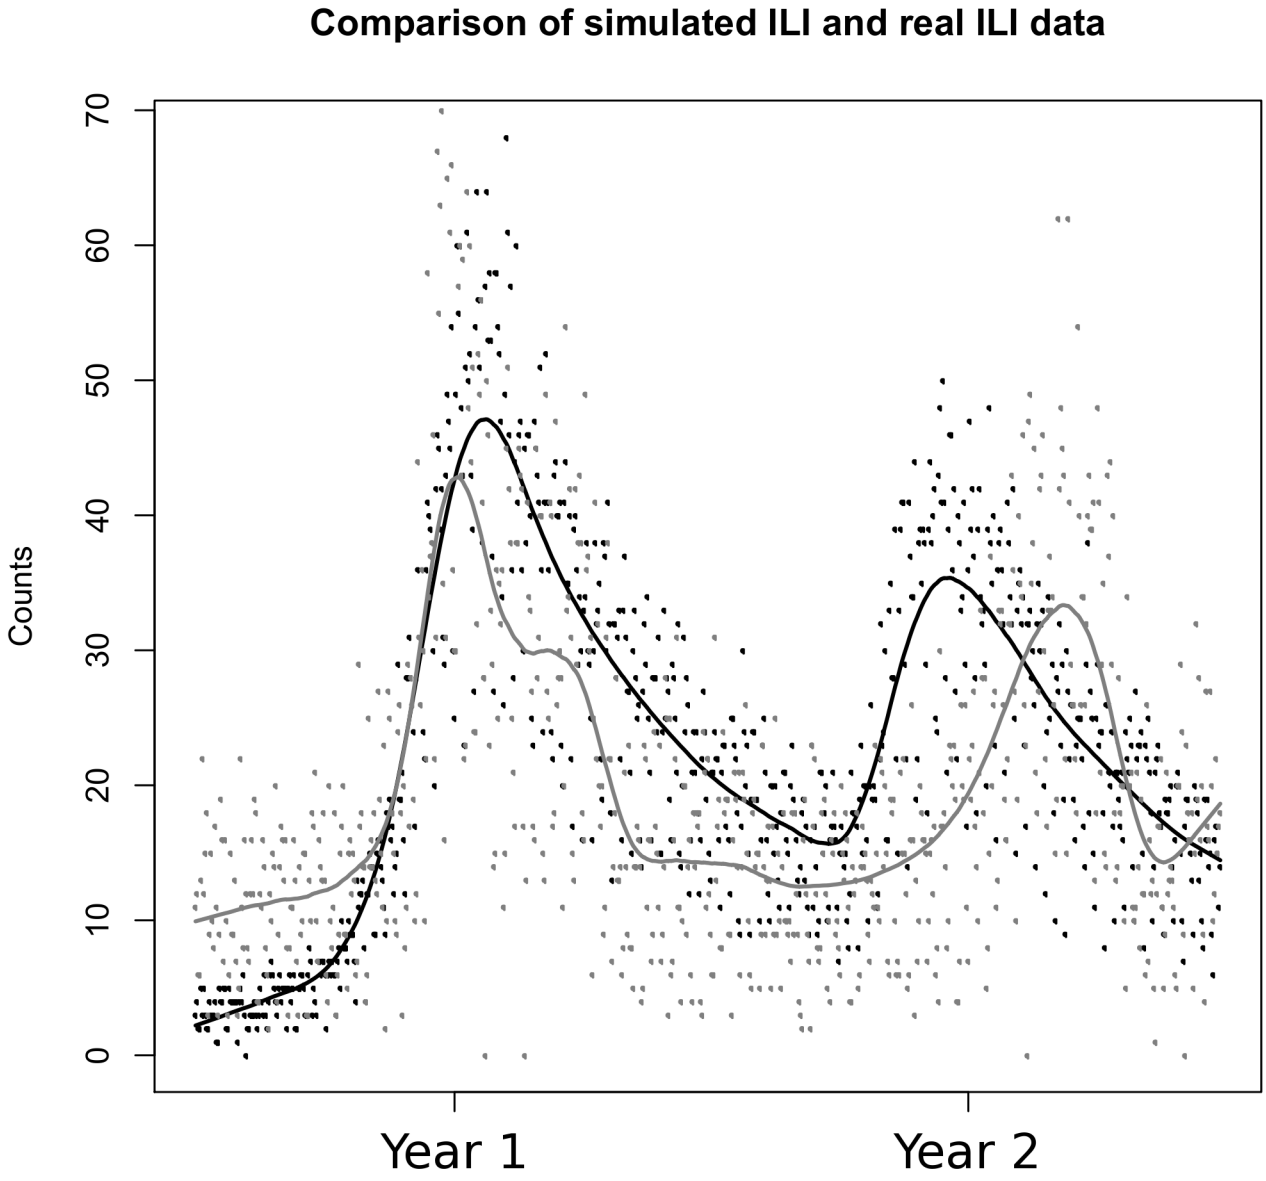

Figure 2. Two years of real ILI data compared to 2 years of data from a single simulation. Each dot represents the total number of cases seen in the surveillance stream on a single day; the curves are a smoothed fit to these data.

Figure 2 shows both real data from HPHC and simulated data. Much as any single influenza season is not the same as previous influenza seasons, the simulated ILI surveillance counts are not the same as the real-world ILI surveillance counts, however they share similar characteristics. ILI surveillance data is characterized by strong seasonal influence and a day of the week pattern, which are evident in the Figure.

The surveillance of the simulated ILI counts without any inserted outbreaks informs decisions about the real-world tenability of a surveillance system. Typically overburdened public health authorities cannot support systems with too many false positives. We found 244 overlapping "clusters" of signals in the 100 years of data in the "Natural 6" system at the 0.001 decision threshold, when the expected number under the null is 36.5. We also saw 547, 669, 394, 626, and 732 in the Natural 12, Natural 18, Uniform 6, 12, and 18 systems respectively.

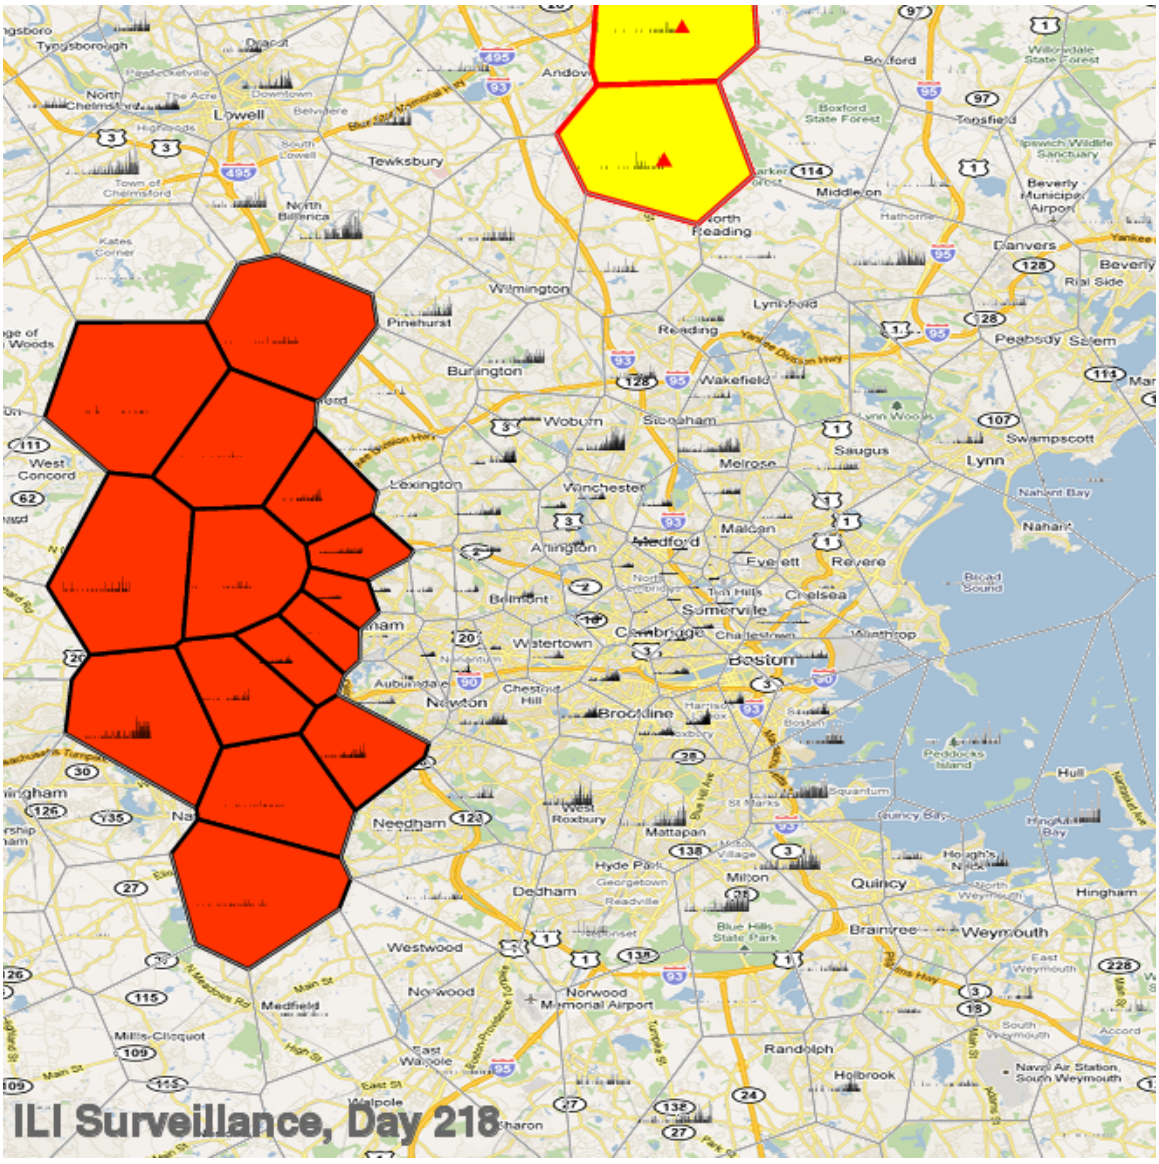

Figure 3. Simulated ILI surveillance data for downtown Boston as captured by the "Natural 6" surveillance system. Surveillance counts per day centered in each zip code location are shown as histograms within each zip code. Detection of an inserted test outbreak (red triangle) is indicated by red bordered zip codes and a false-positive outbreak by black bordered zip codes.

A detailed view of the simulated surveillance data is seen in Figure 3. The zip code boundaries are illustrated as a Voronoi tessellation (23) based on the centroid of the zip code since these are the locations used by the scan-statistic for clustering purposes and not all zip codes are included in the surveillance system. The daily counts of reported cases per zip code are shown as histograms inside the zip code boundaries. The relative contributions of cases from each zip code to the surveillance system can be seen as can the variability between the locations at similar times during the surveillance period. The detection of an inserted outbreak is illustrated with the red highlighted border, whereas the false positive detection of a cluster that doesn't overlap with an inserted outbreak is shown with a thick black border. The fill color of the clustered locations reveals the decision threshold at which they were detected. Animated movies showing this process over the course of an entire year can be found in the supplemental material (and http://ndssl.vbi.vt.edu/insilicoSurveillance/)

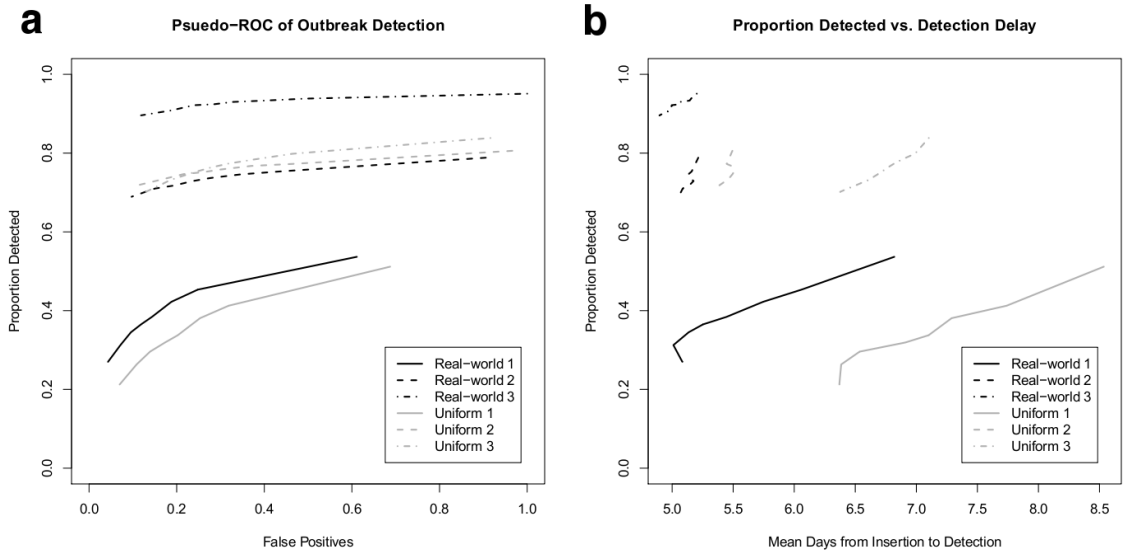

Figure 4. A) Pseudo-ROC curves of outbreak detection. Proportion detected for each surveillance system vs. proportion of all false-positives identified. B) Proportion detected vs. mean days to detection across decision thresholds.

Performance of the six surveillance systems is summarized as pseudo receiver operator characteristic curves (pseudo-ROC) (24) in Figure 4a. True positives are the proportion of inserted outbreaks detected. False positive proportions used the number of false positives detected at the highest decision threshold among all surveillance systems (Natural 18 with p-value = 0.0333) as the denominator for all surveillance systems. As coverage increases so does the likelihood of detection, though the rate of improvement is non-linear and differs in the two coverage distributions. The decision threshold influences the proportion of false-positives more than the probability of detection, though the effect on probability of detection is stronger in the region of decision thresholds likely to have tenable false positive rates (p-value = 0.005 to 0.001). Across all the surveillance systems, sensitivity improves slightly with increased decision threshold

size, however the proportion of false positives increases much more rapidly. Overall, the "Natural" surveillance geography outperforms the "Uniform" surveillance systems. For all detected outbreaks, the delay from insertion to detection can also be measured. The performance of each surveillance system with respect to the mean detection delay and the probability of detection is shown in Figure 4b. Given the experimental design, these delays are the theoretical earliest point of detection by a public health official. The surveillance systems based on the "Natural" surveillance geography consistently show earlier detection than the "Uniform" surveillance geography, even when the probability for detection is essentially the same between the two. Interestingly, higher coverage does not lead to significantly earlier detection: in the case of the "Uniform" surveillance geography, the highest coverage level leads to later detection.

CONCLUSION

Agent-based simulation models incorporating important processes that lead to realistic simulated surveillance data is sufficient for evaluating the performance of a surveillance system for outbreak detection. We demonstrated the use of synthetic methods by conducting an experiment to assess how the level of coverage and geographic distribution in a surveillance system affect performance. This kind of experiment can be used to design, as well as to evaluate, surveillance systems.

Description of the study design and the details of its implementation illustrate the flexibility of the methodology. For example, the analysis could be extended to include additional types of inserted outbreaks, and the model could be altered to supply gastrointestinal surveillance. Keeping the scale and level of detail in the model similar to that of the real-world surveillance system facilitates the analysis of the results and makes the simulation results more compatible with existing evaluation tools and methods. These characteristics support the use of agent-based modeling approaches to a wider variety of public health problems.

Our study shows that the geographic distribution of a surveillance system can have a stronger influence on its ability to detect outbreaks than the level of coverage. This and similar studies can also give guidance to making operational decisions, such as selecting a decision threshold for defining a signal that balances an acceptable number of false-positives against a desired probability of outbreak detection.

The 6-fold difference between the nominal false positive rate of 36.5 signals per hundred years at p-value 0.001 and the observed count of 244 signals in the simulated “Natural 6” data with no inserted outbreaks suggests that the spatio-temporal clustering in the realized simulations is indeed quite different from the null hypothesis of case counts proportional to the population. This confirms that the agent-based model has the desired effect on the pattern of cases. It also shows how the process can help public health authorities anticipate the false positive signaling of such systems.

Efforts to enhance existing surveillance systems can be guided by testing different surveillance systems using methods similar to those described here. For example, the impact of adding a new clinic that would draw patients from additional locations in the population could be simulated to determine the value of the additional information. To optimize a surveillance system's ability to catch outbreaks, a study design that tests the

in silico surveillance system against a variety of simulated outbreaks and outbreak detection algorithms could be conducted. An assessment of the benefit of finer-grain information could also be conducted: rather than base the outbreak detection on the centroids of the home zip code locations, one could add the zip code location of the school or workplace, or replace zip code centroids with anonymized home addresses. The high level of detail is a great benefit of agent-based methods; however, it is also produces what is perhaps their biggest complication. With so many parameters, proper calibration can be time consuming and the opportunity for errors and biases can increase. Similarly, it is sometimes difficult to find suitable data for parameterizing high fidelity models of the real world. For such reasons, these methods are not always intended to provide definitive quantitative predictions; rather, they are intended as a tool for comparisons across several designs or outbreak detection methods. Additionally, the methods are computationally intensive and very few health departments have the resources or personnel to perform them. Fortunately advances in scalable high-performance computing and web services have already begun to provide these resources at relatively low cost and with increasing user-friendliness. Work on developing flexible cyber-infrastructures should make it possible to provide a relatively simple web-based interface for users to conduct studies such as this one in the near future.

Further study using these methods is warranted and will include broader investigation into evaluating different outbreak detection algorithms and surveillance systems. These methods will be useful for developing the next generation of outbreak detection tools. With a very large library of simulated background disease, one can use classification techniques and machine learning to develop "proactive" surveillance systems. These would use state assessments to guide public health officials about where to look for further evidence of an outbreak. For instance, an outbreak that might not create a signal for another week could be preceded by weak non-signaling outbreaks in other regions; when taken individually these might be missed, but when considered as a whole this pattern could improve confidence that a significant outbreak is about to occur. This could reduce the burden on PH departments from following up "unlikely" events and improve PH response times to outbreaks by several days. The techniques needed to accomplish this would be very difficult to apply to limited historical data, thus requiring an approach like this one to generate large volumes of plausible high-fidelity data.

Highly detailed simulation models of infectious disease transmission can be configured for many purposes serving public health. We have demonstrated a flexible framework for using such a model for the evaluation and design of surveillance systems and outbreak detection. While there are limitations to the accuracy with which these models can represent the real world, they can provide sufficiently realistic data at a level of detail that enables previously impossible public health research.

References

1. Buehler JW, Hopkins RS, Overhage JM, Sosin DM, Tong V, Group CW. Framework for evaluating public health surveillance systems for early detection of outbreaks: recommendations from the CDC Working Group. MMWR Recomm Rep 2004;53(RR-5):1-11.

2. Buckeridge DL, Switzer P, Owens D, Siegrist D, Pavlin J, Musen M. An evaluation model for syndromic surveillance: assessing the performance of a temporal algorithm. MMWR Morb Mortal Wkly Rep 2005;54 Suppl:109-15.
3. Kleinman KP, Abrams A, Mandl K, Platt R. Simulation for assessing statistical methods of biologic terrorism surveillance. MMWR Morb Mortal Wkly Rep 2005;54 Suppl:101-8.
4. Eubank S, Guclu H, Kumar VSA, Marathe MV, Srinivasan A, Toroczkai Z, et al. Modelling disease outbreaks in realistic urban social networks. Nature 2004;429(6988):180-4.
5. Halloran ME, Ferguson NM, Eubank S, Longini IM, Cummings DAT, Lewis B, et al. Modeling targeted layered containment of an influenza pandemic in the United States. Proc Natl Acad Sci USA 2008;105(12):4639-44.
6. Koopman JS. Modeling infection transmission- the pursuit of complexities that matter. Epidemiology 2002;13(6):622-4.
7. Lee BY, Brown ST, Cooley PC, Zimmerman RK, Wheaton WD, Zimmer SM, et al. A computer simulation of employee vaccination to mitigate an influenza epidemic. Am J Prev Med 2010;38(3):247-57.
8. Longini IM, Nizam A, Xu S, Ungchusak K, Hanshaoworakul W, Cummings DAT, et al. Containing pandemic influenza at the source. Science 2005;309(5737):1083-7.
9. Parks AL, Walker B, Pettey W, Benuzillo J, Gesteland P, Grant J, et al. Interactive agent based modeling of public health decision-making. AMIA Annual Symposium proceedings / AMIA Symposium AMIA Symposium 2009;2009:504-8.
10. Barrett C, Eubank S, Marathe M. Modeling and Simulation of Large Biological, Information and Socio-Technical Systems: An Interaction Based Approach.; 2006.
11. National Household Travel Survey; 2001.
12. Barrett CL, Beckman RJ, Khan M, Kumar VA, Marathe MV, Stretz PE, et al. Generation and Analysis of Large Synthetic Social Contact Networks. Proceedings of the 2009 Winter Simulation Conference 2009:1003-1014.
13. Atkins K, Barrett CL, Beckman R, Bisset K, Chen J, Eubank S, et al. An Interaction Based Composable Architecture for Building Scalable Models of Large Social, Biological, Information and Technical Systems. CTWatch Quarterly 2008;4(1).
14. Barrett C, Bisset K, Eubank S, Feng X. EpiSimdemics: an efficient algorithm for simulating the spread of infectious disease over large realistic social networks. Proceedings of the 2008 ACM/IEEE conference on Supercomputing 2008.
15. Falsey AR, Hennessey PA, Formica MA, Cox C, Walsh EE. Respiratory syncytial virus infection in elderly and high-risk adults. N Engl J Med 2005;352(17):1749-59.
16. Monto AS. Epidemiology of viral respiratory infections. Am J Med 2002;112 Suppl 6A:4S-12S.
17. Wallace LA, Collins TC, Douglas JDM, McIntyre S, Millar J, Carman WF. Virological surveillance of influenza-like illness in the community using PCR and serology. J Clin Virol 2004;31(1):40-5.
18. CDC. The Flu Season. http://www.cdc.gov/flu/about/season/flu-season.htm. accessed: January 13, 2010

19. Sessa A, Costa B, Bamfi F, Bettoncelli G, D'Ambrosio G. The incidence, natural history and associated outcomes of influenza-like illness and clinical influenza in Italy. Fam Pract 2001;18(6):629-34.
20. Kulldorff M, Heffernan R, Hartman J, Assunção R, Mostashari F. A space-time permutation scan statistic for disease outbreak detection. PLoS Med 2005;2(3):e59.
21. Kulldorff M, Nagarwalla N. Spatial disease clusters: detection and inference. Stat Med 1995;14(8):799-810.
22. Barrett C, Bisset K, Leidig J, Marathe A, Marathe M. An integrated modeling environment to study the co-evolution of networks, individual behavior and epidemics. AI Magazine 2010.
23. Lee… D. Two algorithms for constructing a Delaunay triangulation. International Journal of Parallel Programming 1980.
24. Kleinman KP, Abrams AM. Assessing surveillance using sensitivity, specificity and timeliness. Stat Methods Med Res 2006;15(5):445-64.

Chapter III. *in silico* Instruction: Simulation-Assisted Training and Evaluation for Public Health Decision-Making

Simulation-Assisted Evaluation and Training of Public Health Decision Makers

Bryan Lewis, Stephen Eubank, Chris Barrett, Madhav Marathe, Keith Bissett, James Koopman, Yong Yang, Matthew Samore, Yarden Linvat, and Warren Pettey

ABSTRACT

Background: Simulation-based training and evaluation are frequently employed by a variety of disciplines that require practitioners to make complex decisions, yet this practice has not been widely adopted in public health. A primary obstacle lies in the technical challenge of realistically representing the variety and complexity of a population's health.

Methods: An agent-based model of pertussis transmission and health department contact investigation was developed based on a highly detailed synthetic population of Utah. The simulation was iteratively calibrated and refined to provide sufficiently realistic detail. Scenarios addressing a relevant challenge to pertussis control were developed and presented through software developed to enable users to investigate epidemiologic details to a configurable level of detail.

Results: Twenty unique iterations of pertussis transmission over the entire state of Utah, under nine different combinations of policies were performed. The data produced by the agent-based model were sufficiently detailed to enable the study of the decision making process of public health practitioners. The simulation results were essential to the development of a sophisticated software tool (Epinome) for conducting these studies as well as training public health decision makers. The development process led to important insights into the importance of both the structure of the simulation as well as the method of data presentation.

Discussion: These findings suggest an important role for highly detailed agent-based simulations in improving public health practice. Mature agent-based simulations, advances in high-performance computing, and improved understanding of the mechanisms driving health distributions, enable the construction of environments for training and evaluating public health decision-making.

INTRODUCTION

Rigorous scientific assessment of decision-making is challenging for even the simplest of tasks. When trying to analyze the processes involved in public health

decision-making these difficulties are exacerbated. Agent-based models have reached a level of sophistication that can create an environment that is easily controlled, highly configurable, and sufficiently realistic to permit rigorous studies of how public health practitioners make decisions. While models of this type are often used more directly to inform policy decisions, we focus this study on features of the model and the way data is presented to support an understanding of the decision-making process, without bias as to what the decisions might be.

Simulation-based training and evaluation are frequently employed by a variety of disciplines that require practitioners to make complex decisions [1-3]. This methodology has yet to be widely adopted by the public health field, in part due to the difficulties presented when simulating a sufficiently realistic virtual world. Maturation of agent-based simulations, advances in high-performance computing, and improved understanding of the mechanisms driving health distributions, have created an opportunity to successfully embed public health practitioners in such an environment.

An approach that supports randomized controlled trials assessing the use of decision support tools in public health decision-making is important for the improvement of public health practice. For evaluation purposes, a study could place a decision maker in an environment where the study designer is omniscient and can control the data that is made available, observe how the data is used, record the decisions made, and assess how that decision ranks against other possibilities. This level of configurability also allows for study designs that assess how different decision makers of different levels of experience behave under identical conditions as well as under well-controlled conditions (i.e. having access to richer levels of data or additional tools). For training purposes, a similar environment could be used to challenge the decision maker with a variety of novel situations and through repetitive exposure to these realistic and unique situations, rapidly increase the experience of the public health practitioner.

However, to realize the potential of this approach, it is important that the model outcomes are sufficiently realistic that the subjects respond naturally to them, or at least willingly suspend disbelief. That is, the outcomes must exhibit what is often called ``face validity'' [4]. This is different from other notions of validity, e.g. predictive validity or structural validity. Because the decision-making assessment experiments present extremely detailed pictures of model outcomes, often in ways never expected or calibrated by the modelers, maintaining face validity in these experiments imposes surprisingly stringent requirements on a model. We examine the process of meeting these requirements in the case of a model for pertussis.

A fundamental challenge facing public health practitioners is maintaining situational awareness, i.e. understanding the distribution of the disease in the population. Surveillance systems and epidemiologic studies are employed for this purpose, however, they can be limited in their ability to discern the true incidence of disease. This is especially true for diseases, like pertussis, with high levels of asymptomatic infection and age-specific severity that can complicate diagnosis. In this environment public health practice becomes more challenging, increasing the sensitivity of decision-making to the practitioner's situational awareness. Recreating an environment where the precise distribution of disease is known as well as the as

practitioner's assessment of that distribution gives study designers the ability to study the decision-making process in depth.

To assess the benefits of this approach, technologies needed to study the decision-making process surrounding contact-tracing policies for pertussis control were developed. Recent studies have highlighted the large reservoir of undetected pertussis infections [5-7] as well as the complex role played by waning immunity from vaccination and natural infection[8]. Pertussis is rarely diagnosed amongst adolescents and adults and immunity tends to wane significantly enough by adolescence that outbreaks are frequently reported (need citation) in high schools. To evaluate public health practitioners decision making process they were tasked with identifying outbreaks of pertussis in high schools and deciding what level of intensity should be used at a state-wide level for tracing their contacts.

The methods used to construct the model supporting these environments are described as well as the process used to develop data of sufficient realism to seamlessly immerse the public health practitioners being studied. The importance of different structural features of the data will be discussed as it relates to their relevance for studying different aspects of the public health decision-making process. Several study designs that exploit these aspects are also discussed, as are the qualitative findings of a pilot field experiment.

METHODS

CALIBRATION

Study Population

The methods used to generate a dynamic and highly detailed synthetic population have been described in detail previously[9, 10]. The methods use real world data sources and behavioral models to generate a dataset containing the location of each individual in the population every day. Census data and household location data are used to assign households of individuals drawn from census-block level information to home locations within the geographic area of the census-block. Activity profiles drawn from large activity surveys are assigned to each individual in the household based on household and individual demographics, ensuring that individuals perform activities appropriate to their age and correlated with the structure of the household. A mobility model is used to determine when and where individuals go for each their activities [10].

This proto-population is then updated with additional information specific to the context of the scenario being studied. For instance, the Utah population produced for this study was updated with current school locations and enrollment numbers, for enhanced face validity. Once refined the proto-population is used by a simulation engine to calculate disease transmission, and provide support for dynamic application of the interventions required by the scenario. For this study a distributed discrete event simulation engine, EpiSimdemics[11], was used.

Disease Transmission

Interactions between individuals in the population provide the opportunity for disease transmission. EpiSimdemics simulates the spread of disease through the population based on contacts between uninfected and an infectious individuals [11, 12]. For these simulations, pertussis transmission is a function of the duration of contact, severity of the symptoms of the infectious individual, treatments either individual may have taken, and level of susceptibility of the uninfected individual. The model of pertussis (Figure 1) was based on that developed by Hethcote et al. [13, 14] which allows for 6 levels of limited susceptibility based on the number vaccinations and naturally acquired immunity through infection and the time since the immunity was gained. Additionally, asymptomatic cases of pertussis were only 25% as infectious as a typical case of pertussis, and those treated with antibiotics were 50% less infectious.

To determine an appropriate distribution of susceptibility to infection in the population a compartmental model was developed that incorporated the contact matrix from the agent-based model, the Hethcote disease model, and historical levels of vaccination. The compartmental model was initialized with the pre-vaccine population of Utah and run to 2008 (Yong & Koopman, 2008). The agent-based model used the resulting age-specific distribution as the initial condition for the population-level disease states. The agent-based model was then calibrated to an incidence rate of roughly 490 cases per 100,000 person-years, guided by the large number of undetected pertussis cases reported in the APERT study[7] and Minnesota[6]. Under these conditions (and existing pertussis control policies) the level of infection in the population was found to be stable.

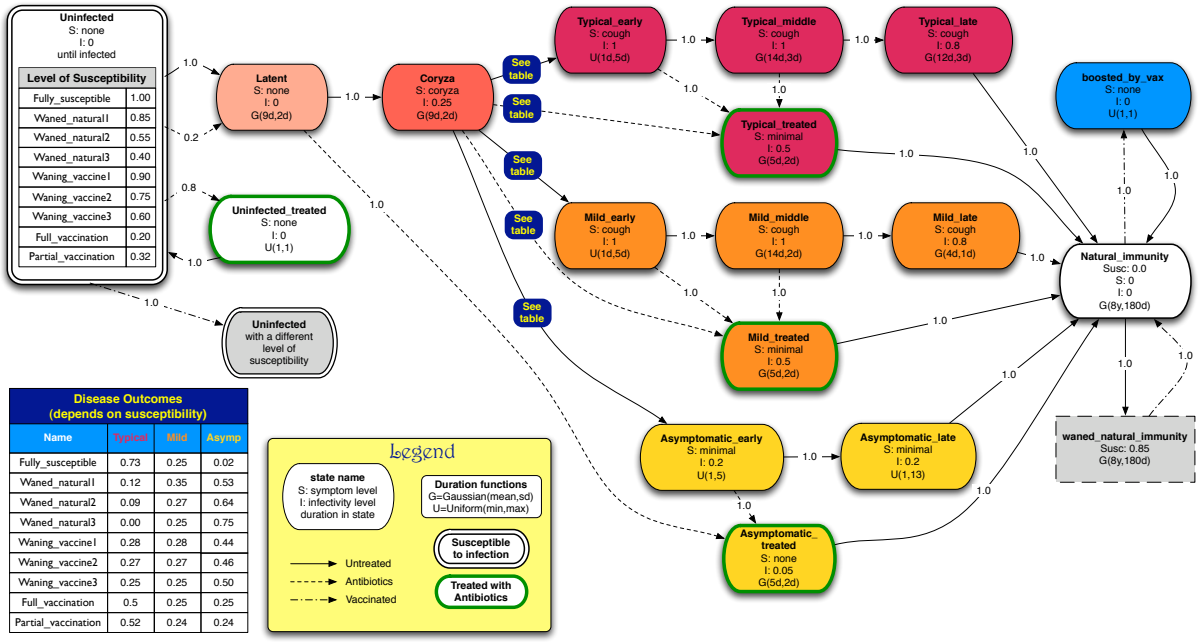

Figure 1. A finite state machine representation of the pertussis disease model, all states, dwell times, transition probabilities, relative level of susceptibility and infectivity, are explicitly represented.

Public Health

Simulated pertussis infections become reported public health cases through several routes. Without any interventions or policies in place the simulation represented the public health reporting system by selecting a fraction of infections as a reported case. The proportion identified as well as the time from symptom onset to case identification was based on the agent's age and the severity of symptoms. These distributions were inferred from Utah Department of Health data. The Utah public health department's policies for identified cases of pertussis were replicated in the model (Table 1). This policy captured additional cases through contact tracing of a subset of the identified cases as well as prevented ongoing transmission through treatment, quarantine, and prophylaxis. With this policy enacted, the level of reported cases was then calibrated to the level of cases observed in the state of Utah.

Interventions are applied dynamically within the simulation based on characteristics of the agent and the policies in place. When cases were identified their contacts were traced based on policies representing three levels of intensity. The least intense policy took no actions on contacts. The moderate level, based on the Utah public health department policies, called for action only on high-risk contacts. The most intense policy stipulated that all contacts be pursued. Since the agent-based model has complete knowledge, all individuals that were co-located with another individual were eligible to be contacts, however, only those contacted for more than an hour and disregarding contacts in shopping centers were counted as contacts. Contacts were classified as high risk if they were less than 1 year old or if they had some other condition that put them at high risk for adverse morbidity with pertussis infection (pregnancy, immuno-compromised, about 10% of contacts). Under immunized contacts (a vaccination history of less than the recommended 4 doses) were also considered high risk. When an individual was identified as a reported case, their contacts were traced and interventions were applied in the following days based on the policies in place. If contact tracing revealed a contact who was symptomatic with pertussis infection, they were identified as case, and the treated as such (i.e. a separate round of contact tracing was performed according the current policy).

#	Policy	Contacts included in intervention	Actions performed on intervened contacts
1	Least intense	None	None
2	Status quo – Maintain current Utah state policy	High-risk contacts only: - Less than 1 year - Under-immunized (incomplete course of vaccines) - High-risk co-morbidity (pregnancy, immuno-compromised, etc.)	- Report all symptomatic cases - Prophylax all uninfected - Quarantine children under 8 years old - Immunize any under-immunized and uninfected individuals

3	Most intense	All contacts of more than an hour duration	Same as policy #2

Table 1. Statewide policy choices for pertussis control, in terms of a package of interventions and their targets.

DATA PRESENTATION

The statewide simulation of pertussis transmission required and produced large volumes of data. These data elements were organized in a relational database for efficient retrieval and storage. Decisions were carefully made as to which data elements to present to the study participant and how to present them.

Data Elements

Structural elements are the core information stored in the population and the simulation structure that governs the operation of the simulation (both directly and indirectly). The simulation engine logs these elements, including all characteristics of the agent, the timing of progression through disease states, and the public health department's interaction with the agent. This data is frequently overly precise (state transitions to the second) and were adjusted to an appropriate level of precision (to the day).

Accessory elements are data elements that are not essential to the simulation itself but are necessary to ensure an immersive experience for the user and maintain face validity. Artificial but realistic names were assigned to all individuals, to permit easier recollection of families and individuals as the user interacted with the data. Similarly, all schools were given their real-world name. These data were added following simulation in a post-processing step.

Interactive Environment Development

A sophisticated interactive environment, Epinome, was created to support the interaction with the data and provide a platform for assessing the decision making process. The concept of Epinome have been previously described[15]. In practice, Epinome provides a user-centric workspace to interact with highly detailed public health information. Since interactive simulations of this complexity are challenging to orchestrate and control, all of the data has been pre-computed. To place the practitioner in a familiar context, the data is revealed in chronological order, day by day. The user can control this process with standard play, pause, fast-forward, and rate controls. The workspace is highly configurable, but is initialized with a variety of views. All cases can be viewed geographically, aggregated into 2-dimensional plots, or as a descriptive line in a table. The individual agent being represented links all these data allowing for all parallel views of the data to be simultaneously highlighted as the user clicks on the different representations.

Epinome, supports task driven information search strategies by providing customizable tools, for example, if the task is finding school based outbreaks, a tool can be produced to aggregate cases by school attendance. Interactive policy decisions are supported by providing cues for decisions at fixed points in simulation playback, and

then proceeding with the appropriate set of simulation data. The study of the decision making process is enabled by supporting individual participant access, logging of all interactions with Epinome, and support for automated survey feedback at the end of a session.

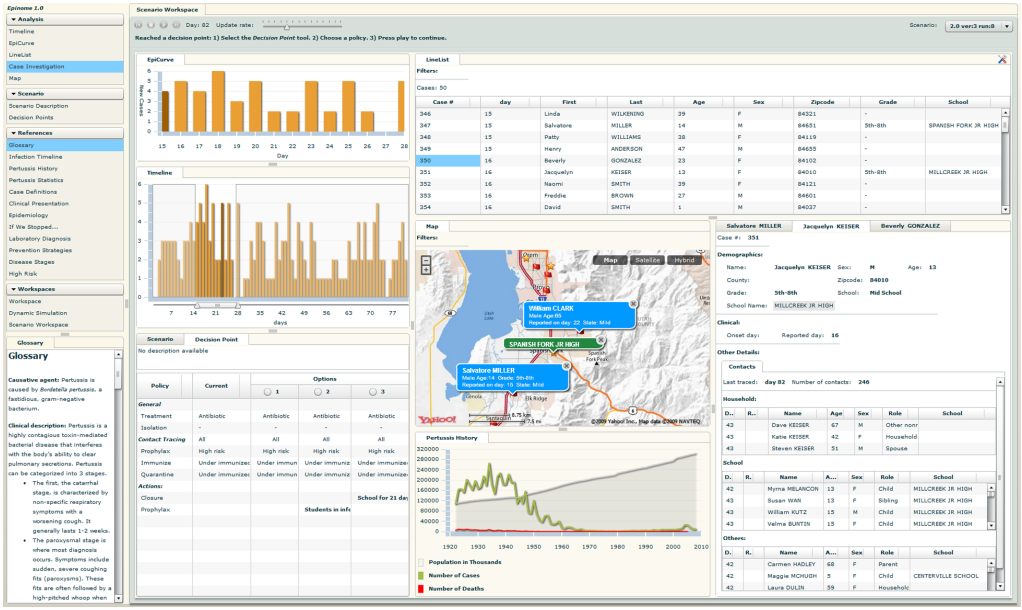

Figure 2. Screenshot of Epinome representing pertussis disease transmission in Utah.

SCENARIO DEVELOPMENT

To challenge the public health decision-making process users were tasked with assessing the state of pertussis infection in the state of Utah, identifying an intentionally created outbreak of pertussis in a school, and selecting an appropriate statewide policy based on their experience. To accomplish these goals simulations were configured to run for 90 days under the status quo policy (policy #2), then adopt one of the 3 policies for 30 days, and then adopt one of the 3 policies again for 90 days, thus 9 unique combinations of policies were simulated for 210 days. A high school was chosen at random as the location of an intentionally seeded outbreak around day 60. For each seeded outbreak, three students who attend the school were infected with pertussis and then flagged for identification as a reported public health case. Twenty unique iterations of these 9 combinations were run (with outbreaks at 20 different high schools).

To train or evaluate a public health decision-maker using this scenario, the participant is surveyed prior to using Epinome to establish an experience baseline. The participant then logs into Epinome, which then loads the data for all 9 combinations for one of the iterations. The participant then explores the data for the first 90-day period and tries to identify the location of the outbreak while learning how to use the tool. At the first decision point they are prompted to make a policy decision, and again at the second decision point 30 days later. Once the entire 210 days have elapsed and the decision maker has thoroughly explored the data, they have the option to view the “ground truth” of the simulation data without being limited to only those cases identified as public health cases. Once completed, the participant is surveyed again to assess what was

learned during the session. Additionally, the record of every action taken by the participant, as recorded by Epinome, can then be thoroughly analyzed.

COMPUTATIONAL REQUIREMENTS

Simulations were conducted on a large cluster [IDAC] (112 4-core 2.2 GHz CPUs for an HPL performance of 1.471 TFlops). Running on 160 cores, a single simulation would take ~30 minutes of wall-clock time, and produce ~500 MB of raw output. Post-processing of this raw data was done using perl scripts which created flat files that contained the most relevant information in a format that could be efficiently loaded into databases, totaling 15MB of processed data per run.

Epinome runs as a web client and is compatible with all major browsers. The client communicates with a central server with a database optimized for efficient information transmission over the web. Epinome supports individual log-ins, allowing customization of the particular scenario presented to the participants. The unique log-ins also allows the tracking of the users survey responses, actions, and decisions during the session.

RESULTS

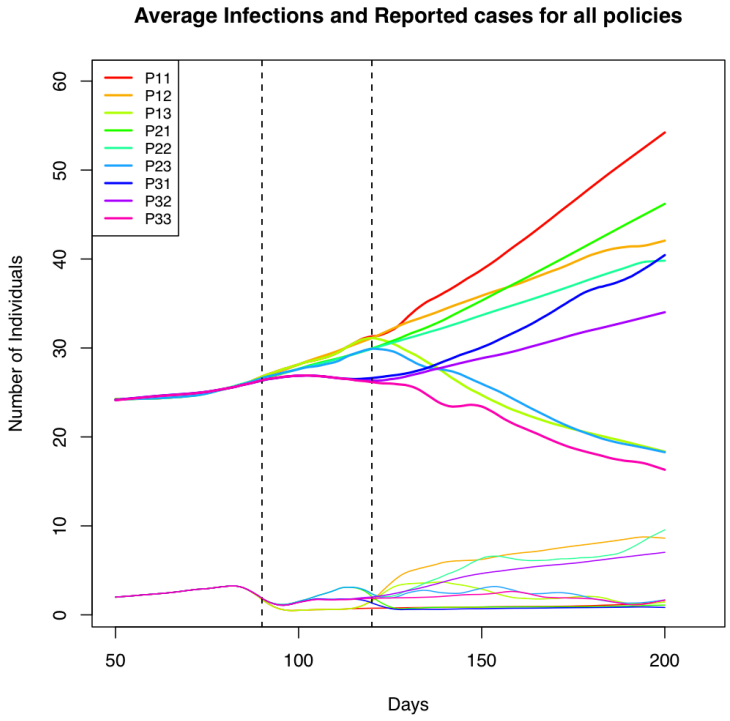

Figure 3. Total daily infections (thick lines) and reported cases (thin lines) for all policy choices averaged over all 20 iterations.

The simulations produced a dynamic social network for the state of Utah with 2.2 million individuals, 241 million unique interactions, and a total of 9.6 million hours of contact per day. The population also contains 442 thousand locations, of which 2536 are schools with a total of 615 thousand students. Figure 3 shows the daily counts of infections and reported cases over the 20 iterations and for each combination of policy choice at the different decision points. Maintaining the least intense policy over the 2 decision points (P11) leads to a rapid rise in the daily counts of infections, however, due to the poor case finding activities a decrease in the reported cases.

The processed simulation data was used to develop a sophisticated software tool (Epinome) for studying the decision making process of public health practitioners as well as support a study of pertussis situational awareness. Figure 2 shows the interface to Epinome and demonstrates the level of detailed information contained in synthetic data. Data are presented in multiple different views allowing simultaneous assessment of the statewide data at using multiple filters to select data and summarizing the data at different levels of aggregation. In Figure 2, the top left shows two epicurves, the timeline on the bottom which shows the history of the epidemiologic data the participant has observed, and the upper epicurve plot which shows the counts in the area of focus (as selected in the historical window). The top right quadrant has a tool that shows a line listing of all cases in the focused area. The line listing displays the demographic details (name, age, gender, school name) of the cases and is sortable by these characteristics. The bottom right quadrant contains the map tool, which places all the cases during the focused period of time in its geographic location. The map tool is interactive, allowing zooming and panning, additionally clicking on the case icons displays a small information blurb about the case. The other tools are specific tools for analyzing the cases that occur in specific schools or other locations. The bottom left quadrant has the decision making tool, that summarized the choices available to the participant as well as what the past decisions were.

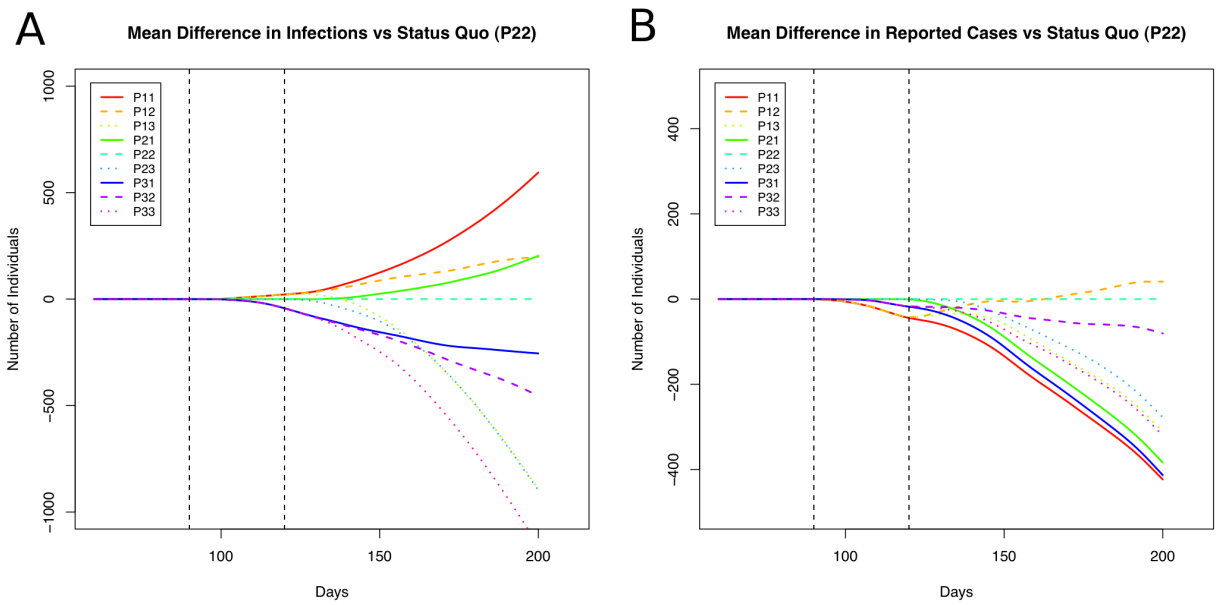

Figure 4. Outcome comparison across policies, cumulative mean differences from the moderate intensity policy (P22). A) The difference in infections (ground truth) B) reported cases

The mean outcomes of the 20 iterations are summarized in Figure 4. The outcomes of each permutation of policies are shown in comparison to maintaining the status quo policy at each decision point (P22, dashed light blue). The actual number of infections (Figure 4A) can be seen to rise most dramatically when the least intense policies are pursued throughout the simulation (P11, red solid line) while the number of reported cases falls the greatest (Figure 4B). The difference between the "ground truth" and the perceived vary not simply in terms of scale but in qualitative behavior. The impact of a short period of intense policies can be seen in comparing the blue solid line vs. the red solid line. While the number of overall reported cases don't vary greatly at day 210. Consistently intense policies (P33) greatly reduce the number of infections and reported cases, but require significantly more activity from the public health department.

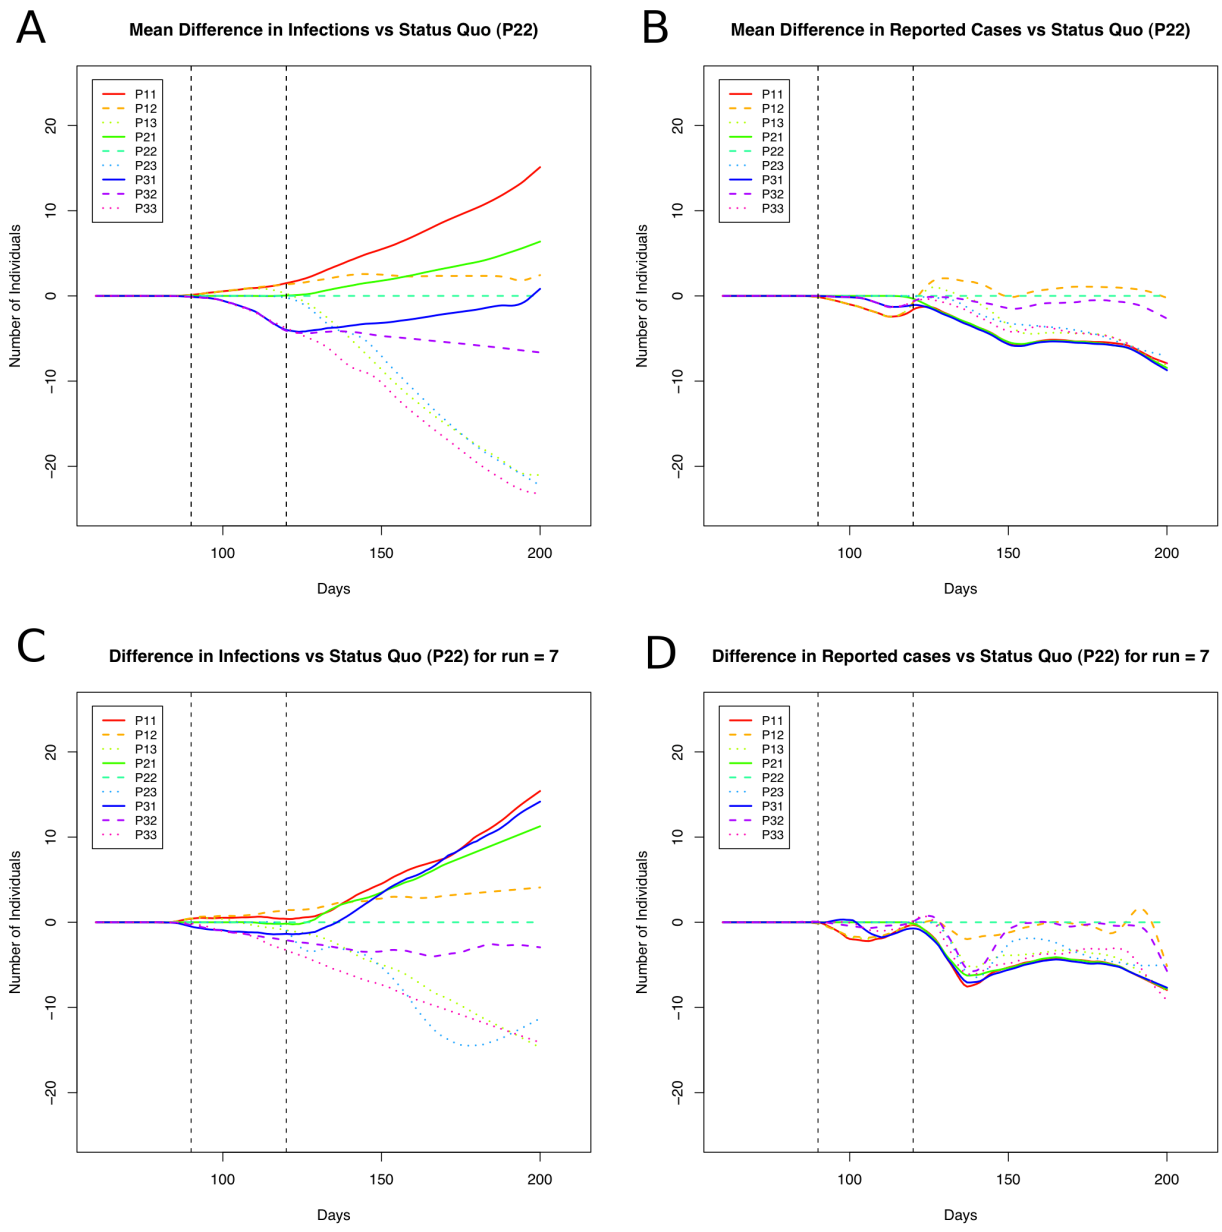

Figure 5. Comparison between average over 20 iterations and an individual stochastic realization (run 7) as a difference from a baseline of the status quo policy (P22). A) Mean infection counts per day B) Mean reported case counts per day C) Single run, infection counts per day D) Single run, reported case counts per day

Mean differences between the policies (Figure 5A & 5B) show the expected behavior for different policy combinations over time. As illustrated in Figure 5, the difference between the actual infections in the population and the perception of those as reflected by reported case counts can behave quite differently. The difference between the expected behavior and any single individual realization (Figure 5C & 5D), however, can vary considerably as well. The group of policy combinations that choose the least

intense intervention as the final decision (solid lines) are well spread out in the expected case (A), however there is minimal difference in the individual case between these three policy combinations. Not only is the scale of the differences between the policies different, but qualitatively the relationships between the policies are flipped in the individual case. The policy combination P13 causes more infections than the combination P12, whereas on average they are around 5 infections per day different, interestingly both policies behave to the similarly to the expected case in terms of the number of reported cases.

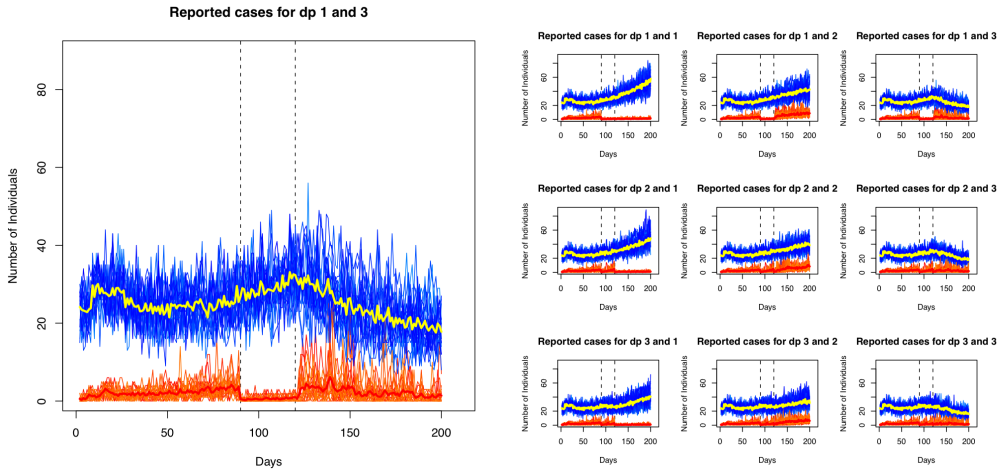

Figure 6. Number of pertussis infections (blue) and number of reported cases (red-orange) per day of simulation. A.) All 20 stochastic realizations for a single set of decisions (trace no contacts at day 90 and trace all contacts at day 120) B.) Same plot but for all 9 combinations of policy decisions.

Figure 6 summarizes the disease transmission and health-department response to pertussis disease transmission represented in the agent-based model. The plots show the number infections and the number of reported cases, and highlight the variability between the different stochastic realizations of the same simulation scenario (Figure 6A). The impact of the different public policy options (Figure 6B) is also demonstrated. Maintaining a high level of stochasticity between the different simulation runs while maintaining a deterministic course of actions is important for comparison between different participants under identical and similar situations. The configurability and control of the scenarios presented provide a degree of control for study designers that can accommodate essentially any design and/or situation to place the study participants.

DISCUSSION

The methods and efforts presented demonstrate an important role for highly detailed agent-based simulations in the field of public health practice. Agent-based simulations of this level of detail are typically used to perform detailed course of action

analysis and used in a "predictive" fashion. We propose a novel use for these simulations as a platform for improving public health practice by enabling the study of the public health decision making process as well as a foundation for developing sophisticated training techniques. Creating data of this configurability, structural integrity, and level of detail would be very difficult to do in a top down manner.

This approach has additional advantages conveyed by the level of detail built into the data. Explicitly representing the many real world objects in the simulation allows for the fusion of many disparate data sources, using the *in silico* data as a foundation. This is particularly relevant in the public health arena when there are often many different sources of data gathered at different levels of aggregation and from different populations by different agencies. The ability to incorporate additional data sources specific to the task and scenario being studied enhances the user's experience. Additionally, detailed representations of the real world increase the compatibility of the simulation results with existing analytic tools and methods that are familiar to the users. These characteristics elicit more natural behaviors from the study participants, improving the applicability of the decision-making observations and the impact of training exercises.

These advantages are only conveyed when the user believes that the data being presented to them is "valid," which requires that the data have both structural and face validity. Extensive knowledge of the processes being represented, attention to the simulation design, and careful calibration of the simulation is essential to ensuring structural validity. Structural validity does not necessarily confer face validity. Subtle differences between the simulated data and the real world, even those that are inconsequential to the operation of the simulation, can lead to a loss to the face validity of the simulated data. Details normally gained through contact investigation in the real world must be available, accurate, and realistic to convey face validity. The timing of case reporting, timing of disease progression, distribution of disease states, and evidence of the impact of interventions must all be available and at an appropriate resolution.

Generating appropriately structured data alone is not sufficient, attention to how the data is presented and the how the user interacts with the data is essential as well. In many ways supporting the development of Epinome is a major finding for these methods. Many of the features designed into it and tools provided would not have been included had lower fidelity and less structured data been used to populate its data structures during the development cycle. Many other disciplines have greatly benefited from enhanced training and evaluation (cite some other examples of simulation aided training: Medicine – Gropper, Design, Management: Sternman) The development of a simulation-aided training tool that can be used to study the public health decision-making process is a necessary tool to add to the public health research armature.

Development of any novel software packages and data management systems to improve public health practice, even those not intended for training and evaluation purposes, can be often be slowed by health information privacy concerns because development is often conducted outside the health department by sub-contractors. Synthetic information produced by these methods would be an excellent way to test out

the functionality as well as the utility of the software being developed. Often the "test data" sets often generated for these purposes are simply placeholders. Not only can real-world levels of variability and volumes of unique data test out the systems being created, but correlations in the data add useful test cases for the software, which can lead to faster implementation in the real world and a more "battle-tested" final product.

Highly detailed data can enable the simulation to achieve sufficient realism, however, ensuring that it conforms to the user's belief in how the system operates imposes some challenges. Including details requires longer periods of data gathering, calibration, and verification. These techniques also require a lot of resources not commonly available to public health departments, both in terms of computational capacity and personnel. However, high-performance computing resources are becoming cheaper and easier to access. Developments in cyber infrastructures hold the potential for advanced systems like the one used in this study to be offered for general use in the coming years.

Development of the scenario presented to the user has important implications. The number of details to calibrate and populate are determined in this process as is the requirements for the simulation itself. At the human-computer interface the scenario dictates how the simulation results are presented to the user and thus the requirements for the software that interfaces with the user. Perhaps most importantly, the goals of the exercise are deeply rooted in scenario development.

The scenario developed for this study was intended to evaluate the decision making process at the policy setting level. A crucial challenge to pertussis control is that immunity from natural infection and vaccination wanes, which creates a large numbers of susceptible individuals in the higher age groups. Infection in these groups is difficult to diagnose, as it is often asymptomatic leading to a reservoir of undetected cases. How aggressively the public health department should pursue identification of these cases depends on the perceived cost-benefit. The scenario requires the user to gain a situation assessment and act on the information. This allows the study designer to evaluate the data gathering process and correlate it with the outcome of the users efforts.

The particular scenario used in this study is a single example of the variety of studies that this data and more generally these methods can support. The level of information resolution is dependent on the details of the disease, the points of intervention under investigation, and the decision making process to be studied. Using this established data and study framework there are several aspects of the decision-making process that could be studied. The sensitivity to the resolution of available information could be studied, by aggregating the simulation results to several different levels. Important aspects of how public health practitioners absorb information could be studied by providing different styles of visualizations or analytic tools to different participants. The sensitivity of the actual decisions reached to decision maker's assumptions at the beginning of the exercise and the information they were able to uncover could be scientifically assessed as well.

Highly detailed and appropriately structured data representing disease spread in a realistic population provide an ideal platform for supporting the study of public health

decision-making. The ability to finely tune the scenarios and at the same time provide highly detailed levels of data provides the necessary foundation for an environment to comprehensively assess how decision makers use information for their decisions as well as test the utility of decision support tools. The construction of a sophisticated interactive environment for conducting these studies as well as the development of the scenario for studying the decision making process for statewide pertussis control illustrates the utility of this framework for studying the decision making process.

REFERENCES

1 Reuter C, Pipek V, Mueller C. Computer Supported Collaborative Training in Crisis Communication Management. Proceedings of the 6th International International Conference on Information Systems for Crisis Response and Management (ISCRAM 2009), Göteburg 2009.

2 Van Nortwick SS, Lendvay TS, Jensen AR, et al. Methodologies for establishing validity in surgical simulation studies. Surgery 2010;**147**:622-30.

3 Wahidi MM, Silvestri GA, Coakley RD, et al. A prospective multicenter study of competency metrics and educational interventions in the learning of bronchoscopy among new pulmonary fellows. Chest 2010;**137**:1040-9.

4 Wright Forrester J. Chapter 13. Judging Model Validity. Industrial dynamics 1961:464.

5 Purdy KW, Hay JW, Botteman MF, et al. Evaluation of strategies for use of acellular pertussis vaccine in adolescents and adults: a cost-benefit analysis. Clin Infect Dis 2004;**39**:20-8.

6 Strebel P, Nordin J, Edwards K, et al. Population-based incidence of pertussis among adolescents and adults, Minnesota, 1995-1996. J Infect Dis 2001;**183**:1353-9.

7 Ward JI, Cherry JD, Chang S-J, et al. Efficacy of an acellular pertussis vaccine among adolescents and adults. N Engl J Med 2005;**353**:1555-63.

8 Wearing HJ, Rohani P. Estimating the duration of pertussis immunity using epidemiological signatures. PLoS Pathog 2009;**5**:e1000647.

9 Eubank S, Guclu H, Kumar VSA, et al. Modelling disease outbreaks in realistic urban social networks. Nature 2004;**429**:180-4.

10 Barrett CL, Beckman RJ, Khan M, et al. Generation and Analysis of Large Synthetic Social Contact Networks. Proceedings of the 2009 Winter Simulation Conference 2009:1003-14.

11 Barrett C, Bisset K, Eubank S, et al. EpiSimdemics: an efficient algorithm for simulating the spread of infectious disease over large realistic social networks. Proceedings of the 2008 ACM/IEEE conference on Supercomputing 2008.

12 Atkins K, Barrett CL, Beckman R, et al. An Interaction Based Composable Architecture for Building Scalable Models of Large Social, Biological, Information and Technical Systems. CTWatch quarterly : cyberinfrastructure technology watch 2008;**4**:46-53.

13 Coudeville L, van Rie A, Andre P. Adult pertussis vaccination strategies and their impact on pertussis in the United States: evaluation of routine and targeted (cocoon) strategies. Epidemiol Infect 2008;**136**:604-20.
14 Van Rie A, Hethcote HW. Adolescent and adult pertussis vaccination: computer simulations of five new strategies. Vaccine 2004;**22**:3154-65.
15 Parks AL, Walker B, Pettey W, et al. Interactive agent based modeling of public health decision-making. AMIA Annual Symposium proceedings / AMIA Symposium AMIA Symposium 2009;**2009**:504-8.

Conclusions

The work presented in this dissertation demonstrates the essential role of highly detailed simulations in support of Public Health decision-making. The breadth of work includes directly supporting decision makers as well as the development of novel tools to enable public health research. Some of the approaches used in this dissertation are extensions of previously established domains in public health, while others are relatively novel and point to new frontiers in computational epidemiology.

As the role of simulations in support of Public Health has evolved from simple mathematical models, to increasingly complicated models that require simulation, to highly complicated agent-based simulations, many challenges have arisen at the interface between modeler and public health decision maker. Many of these challenges were encountered in the course of this study, some were overcome (to varying degrees), and others were not. Frequently problems resulted simply from the complexity of problems and systems being studied. Addressing problems that span multiple domains necessitated collaborations with multiple teams from different institutions, which challenged both communication as well as organization between the teams. Encountering these challenges in the course of this study provided many learning opportunities and built a firm foundation of experience for future efforts.

Continuing this work will present exciting opportunities to make *in silico* Public Health research more widely adopted and understood. Directly supporting public health decision making as detailed in Chapter I and Appendices B, C, and D will continue to evolve. Meeting these needs will require continued software development, changes in public health data management, and novel analytic methods. The work in both Chapters II and III prove the concept that synthetic data generated through simulations can enable novel public health research. This paves the way for the refinement of many existing methodologies, as well as the initiation of research that was previously impossible. The most exciting opportunity lies in the potential of these techniques to train the next generation of public health practitioners as well as improve existing public health practice.

SUMMARY

Chapter I

Directly support public health decision-making mobilized a series of evolving efforts over the course of several years. Initial engagement with decision makers began with the study presented in the manuscript. This effort took nearly a year to complete as there were several groups to coordinate and the technologies used for the simulation were still under development. Following this work an effort to evaluate antiviral “Medkit” distribution plans brought another request for simulations (Appendix B). VBI was ready to respond with a detailed simulation, thanks to the recently developed simulation platform EpiFast, which could execute simulations much more rapidly. This effort took a couple months to complete as it required further extensions of the new simulation platform and the requirements from the decision-makers were not as focused as the

previous request. The results of these studies went into devising the HHS Pandemic Influenza Response Plan.

When the pandemic finally did arrive, decision-makers again asked modelers for assistance in enhancing situational awareness as well as responding to time-sensitive policy decisions. Following the emergence of, what would be identified as, pandemic 2009 influenza A (H1N1) virus there was rush to determine the disease characteristics and their potential impact on US populations (Appendix C). The study took a couple weeks to complete, but with a significant difference in that much of the work was conducted by the government analysts being tasked for answers by the decision-makers. The work presented in Appendix C was done in support of the analysts and explored the range of potential disease parameterizations. What is not shown is roughly 100 simulation sets run by solely by the analysts that were used to support daily briefings over the course of the early emergence of the virus. In the fall of 2009 after the pandemic struck with full strength another study was conducted to assess the likelihood that a "3rd wave" of the epidemic would appear following the end-of-year holiday season (Appendix D). This study took less than a week to support due to additional advances in the simulation software and the intensive interactions between an expert user and the decision-makers. The improbability of a 3rd wave determined by this study was used by DHHS to direct their efforts in the aftermath of the main 2nd wave of the epidemic.

These experiences demonstrate that highly detailed simulations can directly support public health decision-making even for complex scenarios and under time-sensitive conditions. The evolutionary process followed for the development of the tools was instrumental in the increasingly small periods of time needed to fulfill requests, as was the ability for the decision-makers to understand the capabilities of the simulations. Further development and refinement of these tools and continued engagement with decision-makers will lead to a robust suite of computationally based decision making tools for Public Health.

Chapter II

Outbreak detection is a core part of public health practice and is a universally recognized problem that requires significant effort from public health practitioners and researchers alike. Considerable research efforts continually describe innovative approaches to collecting better surveillance data, interpreting the surveillance data, and ways to respond to detected signals. However, very few methods exist that allow one to evaluate these methods side-by-side without conducting large-scale population based studies. The inspiration to apply existing outbreak detection algorithms to simulated surveillance data was sparked during a MIDAS meeting. One of the missions of MIDAS is to foster collaborations between the different groups, under these conditions where collaboration was incentivized the initial concept was able to slowly mature from the simple idea along into the project described in Chapter II. Without close collaboration between two groups with diverse areas of expertise these two methodologies could not have been so successfully merged.

The research in Chapter II demonstrates that highly detailed simulations can generate synthetic information of sufficient resolution to serve as a platform for

evaluating complex methodologies. The particular application to surveillance system design and outbreak detection algorithm evaluation addresses a very relevant public health problem. The establishment of this methodology and its provision as a web service should eventually provide many local health departments a means to improve their currently employed outbreak detection processes.

Chapter III

Public health decision-making is a difficult process to capture, let alone in a repeatable and controlled manner. The research presented in Chapter III was a CDC funded collaboration that successfully constructed a set of tools and a methodology to study the process in detail. An obstacle to using a simulation-based approach was generating an environment that replicated the user's view of the real world with sufficient detail so as to maintain face validity with the participant. To generate the complex structurally sound skeleton needed to maintain face validity with the participants required highly detailed synthetic population data combined with a highly calibrated and realistic disease and behavior simulations. While the selection of a real world disease for the scenarios required a very labor-intensive calibration process, the context of the scenario is familiar for the participants. Following an intensive week of on-site interaction, the synthetic epidemiologic data was processed into a form compatible with the interactive environment developed by collaborators in Utah. The combination of the highly detailed well-structured data and the environment that enabled seamless interaction with this data that ultimately yielded an effective tool (described in Appendix G).

The construction of these tools and methods enables detailed studies of the public health decision-making process, as well as provides the opportunity for a robust training environment. Aristotle extols the benefits of extensive training, "Excellence is an art won by training and habituation." These tools can be used across multiple domains and by public health practitioners with a variety of responsibilities, from those who directly interact with the public to those deciding statewide policies. The ability to embed public health practitioners in a realistic environment and present them with relevant challenges has the potential for greatly improving their ability to make effective decisions. Additionally, the in depth study of the decision-making process can further eliminate obstacles and develop tools to further enable efficient public health decision-making.

LESSONS LEARNED

Much has been learned during the course of conducting the research presented in this dissertation. Many new technical skills have been acquired through the myriad challenges encountered. A familiarity with a variety of new knowledge domains, scientific methods, and the process of scientific research in general has been gained. The nature of this trans-disciplinary research has fostered working relationships with individuals across a wide spectrum of scientific disciplines. In many ways, the

numerous small lessons learned represent the real benefit of completing this course of study, not just the small fraction of effort formally described in the above manuscripts.

On a fundamental level, having a greater number of tools available for tackling a problem can greatly improve research effectiveness. To an increasing extent open-source and freely available software tools are proving to be valuable resources. Programming within the UNIX shells or with Python, Perl, or R all were used in the course of this study, as were several software tools like Google Earth, a variety of specialized R packages, gnuplot, GraphViz, GIMP, OmniGraffle, and many others. As these simulations all required high-performance computing resources fluency with distributed computing concepts and its tools was also acquired.

The numerous public presentations of the progress of these research projects afforded many chances to distill and communicate important ideas, concepts, and findings. Being able to effectively communicate the details of complex projects is nearly as important as the actual execution of the work. One of the more effective ways to communicate complex data or convey a multi-layered concept is through visualization. In the course of this study, several courses in visualization were attended. Significant effort was put into generating visual representations of data and learning skills that enable the generation of high impact visualizations. These visualizations can be equally useful as a tool for the researcher to understand their own data as it often reveals previously undetected patterns or associations that had not yet been discovered.

The projects presented in this dissertation, could not have been completed as effectively without collaborating with other groups. Collaboration presents opportunities for all participants to gain from the strengths of the others, however, this does not always happen. Early communication about the goals of the project, what success will mean, and clear understandings of roles can improve the odds of an effective collaboration. Participation in a less effective collaboration in the course of this study emphasized the importance of communication, early agreement on the specifics of the goals, and establishing a realistic timetable. A more effective collaboration was aided by the adoption of several software tools that facilitated long distance collaboration. The availability of free and high-quality video conferencing (Skype and Google Talk) along with collaborative document sharing (DropBox and Google Docs) has enabled working groups separated by geography to collaborate more efficiently. These tools alone cannot make up for shortfalls in interpersonal organization; it can lower the friction that physical distance can create. Collaboration as an ideal is not sufficient, even if well-organized and armed with effective tools, without a true need or incentive to collaborate, collaborative efforts seem to fail.

Even with effective collaborations a team science approach is necessary for tackling problems as complex as the development of innovative epidemiological simulation software. The experience of working closely with software developers to refine the simulation software reinforced the importance of this approach. Making decisions about key capabilities to add and prioritizing the actual methods implemented must be made with knowledge about epidemiology, infectious disease, and immunology. Similarly devising efficient algorithms and organizing the immense software projects necessary to deliver these sophisticated simulations requires

individuals with a great degree of computer science and software engineering skill. Furthermore, harnessing the large volumes of data needed as a foundation for many of these studies also requires a diverse set of skills. These are all tasks too specific and broad for a single individual to accomplish efficiently and effectively, and require a team science approach.

The multiple opportunities to be involved in efforts that informed federal policy decisions provided some insight into how information and concepts are taken from scientific research and absorbed by the decision making process. An appreciation for the process and the unequal influences on the decisions rendered that can be exerted by the variety of "interests" involved is instructive. These higher-level policy-setting processes are a complex system, much like the systems it seeks to influence. The variability and robustness these processes exhibit are their strength, but also make them frustratingly difficult to understand and predict.

LIMITATIONS

The many advantages of *in silico* science are equally balanced by its many limitations. The complexities of the computational technologies needed to successfully conduct these kinds of studies are an inherent obstacle. Even with this obstacle removed through either computational expertise or simplified software requirements, the process can require considerable time and resources.

While there are methods to improve the value of sparse data, simulations are greatly limited by their intrinsic need for data. The quality of the information generated by simulations is highly dependent on the quality of the data used to parameterize them. In many situations a simulation-based approach is simply not appropriate. To further compound this problem, it is often difficult to assess the utility of simulation results prior to conducting the *in silico* study.

Difficulties associated with the dearth of data extend to efforts to validate and verify the simulation results. Best practices require that data not used in the parameterization of the study be used for these purposes, however, in an environment with limited amounts of relevant data this can be impossible. Alternative approaches to establishing confidence in the model can be used, however, reliance on high quality data remains essential to the integrity of simulation studies.

As a relatively new technology in public health research the strengths and weaknesses of computational approaches are not well understood. This lack of familiarity limits often leads fellow researchers to define the whole approach as a "black box" that they can't understand. This phenomenon impedes efficient communication about the details of the simulation, which leads other public health researchers to either completely trust or distrust the "black box." This ultimately weakens the simulation study as it is deprived of the usual level of healthy skepticism and insightful criticism.

NEXT STEPS

Several of these limitations are addressed by the future plans for these projects. The next steps in the line of work demonstrated in Chapter I and its associated

appendices, is to continue to build on the technologies developed. The engagement with decision-makers to provide simulation results to address current public health problems will continue to provide novel challenges to the simulation techniques and the study designs that are possible. Exploration of diseases beyond influenza and engagement with other agencies is also planned.

The end goal of these advances is to give a more robust environment for the decision makers to assess their options and weigh their choices. In practice this environment would facilitate the interaction with multiple different simulations and many different layers of synthetic data. The potential exists for this environment to accommodate the many actual decision makers that must interact in the real world, interfacing with each other in a virtual world, testing out potential responses to challenging problems they are all facing.

The work discussed in Chapter II was conducted in collaboration with the Harvard Medical School MIDAS group. Currently there are plans to recreate the work presented in the manuscript using multiple different outbreak detection algorithms, in an attempt to systematically assess their benefits and limitations. Additionally a study that attempts to quantify the influence of more sophisticated surveillance system designs is planned. Essentially this study will investigate how the resolution of information collected by the surveillance system can affect its sensitivity.

In the spirit of further enabling previously impossible public health research, there are plans to provide *in silico* surveillance data as a web service. Ideally the users would be able to configure the disease, population, and surveillance system through a web interface and have "the cloud" deliver realistic synthetic surveillance data. Advances in "cloud computing" make it realistic that the computational resources to calculate and serve the synthetic data will be affordable for many researchers at universities and at the state and local health department level. Significant effort will have to go into modifying the current methods to be compatible with "the cloud" however; much of the work in terms of design has already taken place. Provision of these services supplements the publication and other forms of dissemination of the methods. The ability to exactly replicate what was reported is one of the main advantages of *in silico* methods.

The project described in Chapter III established a useful set of tools and a dataset for evaluating public health decision-making. A study was piloted in Utah, and a full-fledged study should be conducted. Unfortunately the funding for the collaborative phase of this study has expired, in the absence of a clear incentive further collaboration is unlikely. However, the project generated a very well calibrated and detailed model of pertussis transmission, which will be used to study pertussis policies and advocate for improvements. The lessons learned and the relationships forged during this project will lead to further support for the interactive decision making environment described above.

While this dissertation makes a compelling case for the essentiality of highly detailed simulations for Public Health decision-making, many obstacles remain before these approaches can be fully utilized. The planned extensions will provide additional engagement with public health practitioners to help ensure the relevance of the individual approaches. However, in the long term, tools like these and a more thorough

treatment of epidemiologic modeling needs to be better integrated into the curriculums of public health programs. More training and exposure of public health practitioners to these technologies will better equip them to understand and interpret the ever-increasing number of simulation-based public health research studies. It is only through engagement from public health practitioners at all levels and across all domains the full potential of simulation-based studies can be realized.

FINAL THOUGHTS

A robust public health system is essential to a healthy and functioning society. The increasing interdependence of the world holds great promise, but also brings many challenges. As recent examples of these threats from the H1N1 pandemic to the rise of international terrorism demonstrate, all-hazards planning is the best strategy for mitigating these threats. Public Health has a crucial role to play in this strategy and needs to draw on all resources at its disposal. The work presented in this dissertation represents initial implementations of applying advanced simulation science coupled with high-performance computing to relevant public health problems. They are just a few examples of the potential an *in silico* Public Health approach holds, and demonstrate its essential role in supporting public health decision-making.

References

1. Blower S, Bernoulli D. An attempt at a new analysis of the mortality caused by smallpox and of the advantages of inoculation to prevent it. 1766. Rev Med Virol. 2004;14(5):275-88.
2. Halloran ME, Ferguson NM, Eubank S, Longini IM, Cummings DAT, Lewis B, et al. Modeling targeted layered containment of an influenza pandemic in the United States. Proc Natl Acad Sci USA. 2008;105(12):4639-44.
3. Dietz K, Heesterbeek JA. Bernoulli was ahead of modern epidemiology. Nature. 2000;408(6812):513-4.
4. Bernoulli D. Essai d'une nouvelle analyse de la mortalite causee par la petite verole. Mem Math Phy Acad Roy Sci Paris. 1766.
5. Kermack WO, McKendrick AG. Contributions to the mathematical theory of epidemics--I. 1927. Bull Math Biol. 1991;53(1-2):33-55.
6. Louz D, Bergmans HE, Loos BP, Hoeben RC. Emergence of viral diseases: mathematical modeling as a tool for infection control, policy and decision making. Critical reviews in microbiology. 2010.
7. Lipsitch M, Cohen T, Cooper B, Robins JM, Ma S, James L, et al. Transmission dynamics and control of severe acute respiratory syndrome. Science. 2003;300(5627):1966-70.
8. Rosenthal A, Mork P, Li MH, Stanford J, Koester D, Reynolds P. Cloud computing: a new business paradigm for biomedical information sharing. J Biomed Inform. 2010;43(2):342-53.
9. Blower SM, McLean AR. Prophylactic vaccines, risk behavior change, and the probability of eradicating HIV in San Francisco. Science. 1994;265(5177):1451-4.
10. Ferguson NM, Donnelly CA, Anderson RM. The foot-and-mouth epidemic in Great Britain: pattern of spread and impact of interventions. Science. 2001;292(5519):1155-60.

11. May RM, Anderson RM. Transmission dynamics of HIV infection. Nature. 1987;326(6109):137-42.
12. Yorke JA, Nathanson N, Pianigiani G, Martin J. Seasonality and the requirements for perpetuation and eradication of viruses in populations. American journal of epidemiology. 1979;109(2):103-23.
13. Atkins K, Barrett CL, Beckman R, Bisset K, Chen J, Eubank S, et al. An Interaction Based Composable Architecture for Building Scalable Models of Large Social, Biological, Information and Technical Systems. CTWatch quarterly : cyberinfrastructure technology watch. 2008;4(1):46-53.
14. Barrett C, Bisset K, Eubank S, Feng X. EpiSimdemics: an efficient algorithm for simulating the spread of infectious disease over large realistic social networks. Proceedings of the 2008 ACM/IEEE conference on Supercomputing. 2008.
15. Wright Forrester J. Chapter 13. Judging Model Validity. Industrial dynamics. 1961:464.
16. Kopec JA, Finès P, Manuel DG, Buckeridge DL, Flanagan WM, Oderkirk J, et al. Validation of population-based disease simulation models: a review of concepts and methods. BMC Public Health. 2010;10:710.
17. Glezen WP, Couch RB. Interpandemic influenza in the Houston area, 1974-76. N Engl J Med. 1978;298(11):587-92.
18. Yasuda H, Yoshizawa N, Kimura M, Shigematsu M, Matsumoto M, Kawachi S, et al. Preparedness for the spread of influenza: prohibition of traffic, school closure, and vaccination of children in the commuter towns of Tokyo. J Urban Health. 2008;85(4):619-35.
19. Control CfD, (CDC) P. Overview of Influenza Surveillance in the United States. CDC webpage. 2010.
20. Reed C, Angulo F, Swerdlow D, Lipsitch M, Meltzer M, Jernigan D, et al. Estimates of the prevalence of pandemic (H1N1) 2009, United States, April–July 2009. Emerg Infect Dis. 2009;15(12):2004-7.
21. Kleinman KP, Abrams A, Mandl K, Platt R. Simulation for assessing statistical methods of biologic terrorism surveillance. MMWR Morb Mortal Wkly Rep. 2005;54 Suppl:101-8.
22. Kulldorff M, Heffernan R, Hartman J, Assunção R, Mostashari F. A space-time permutation scan statistic for disease outbreak detection. PLoS Med. 2005;2(3):e59.
23. McBrien KA, Kleinman KP, Abrams AM, Prosser LA. Use of outcomes to evaluate surveillance systems for bioterrorist attacks. BMC Med Inform Decis Mak. 2010;10:25.
24. Qin M, Wang D-y, Huang F, Nie K, Qu M, Wang M, et al. Detection of pandemic influenza A H1N1 virus by multiplex reverse transcription-PCR with a GeXP analyzer. J Virol Methods. 2010;168(1-2):255-8.
25. Jeong I, Lee C-h, Kim DK, Chung HS, Park SW. Mild form of 2009 H1N1 influenza infection detected by active surveillance: implications for infection control. Am J Infect Control. 2010;38(6):482-5.
26. Banerjee T, Pensi T, Banerjee D. Sensitivity of paediatric AIDS score vs. WHO case classification in Indian children--a retrospective study. J Trop Pediatr. 2009;55(2):91-6.
27. Choi K-w, Wong N-s, Lee L-y, Lee S-s. Surveillance of febrile patients in a district and evaluation of their spatiotemporal associations: a pilot study. BMC Public Health. 2010;10:84.
28. Sugiura H, Ohkusa Y, Akahane M, Sugahara T, Okabe N, Imamura T. Construction of syndromic surveillance using a web-based daily questionnaire for health and its application at the G8 Hokkaido Toyako Summit meeting. Epidemiol Infect. 2010;138(10):1493-502.
29. Carneiro HA, Mylonakis E. Google trends: a web-based tool for real-time surveillance of disease outbreaks. Clin Infect Dis. 2009;49(10):1557-64.

30. Eysenbach G. Infodemiology and infoveillance: framework for an emerging set of public health informatics methods to analyze search, communication and publication behavior on the Internet. J Med Internet Res. 2009;11(1):e11.
31. Regan JF, Makarewicz AJ, Hindson BJ, Metz TR, Gutierrez DM, Corzett TH, et al. Environmental monitoring for biological threat agents using the autonomous pathogen detection system with multiplexed polymerase chain reaction. Anal Chem. 2008;80(19):7422-9.
32. Buckeridge DL, Switzer P, Owens D, Siegrist D, Pavlin J, Musen M. An evaluation model for syndromic surveillance: assessing the performance of a temporal algorithm. MMWR Morb Mortal Wkly Rep. 2005;54 Suppl:109-15.
33. Buehler JW, Hopkins RS, Overhage JM, Sosin DM, Tong V, Group CW. Framework for evaluating public health surveillance systems for early detection of outbreaks: recommendations from the CDC Working Group. MMWR Recomm Rep. 2004;53(RR-5):1-11.
34. Nair M, Webster P. Education for health professionals in the emerging market economies: a literature review. Med Educ. 2010;44(9):856-63.
35. Smith RD. The application of information technology in the teaching of veterinary epidemiology and public health. J Vet Med Educ. 2004;30(4):344-50.
36. Harrity S, Jackson M, Hoffman H, Catanzaro A. The National Tuberculosis Curriculum Consortium: a model of multi-disciplinary educational collaboration. Int J Tuberc Lung Dis. 2007;11(3):270-4.
37. Tyler IV, Hau M, Buxton JA, Elliott LJ, Harvey BJ, Hockin JC, et al. Canadian medical students' perceptions of public health education in the undergraduate medical curriculum. Acad Med. 2009;84(9):1307-12.
38. Griffith JM, Sorenson JR, Bowling JM, Jennings-Grant T. Assessment of an interactive computer-based patient prenatal genetic screening and testing education tool. Health Educ Behav. 2005;32(5):613-26.
39. Wisner KL, Logsdon MC, Shanahan BR. Web-based education for postpartum depression: conceptual development and impact. Arch Womens Ment Health. 2008;11(5-6):377-85.
40. Glasgow RE. Interactive media for diabetes self-management: issues in maximizing public health impact. Med Decis Making. 2010;30(6):745-58.
41. Pederson LL, Blumenthal DS, Dever A, McGrady G. A web-based smoking cessation and prevention curriculum for medical students: why, how, what, and what next. Drug Alcohol Rev. 2006;25(1):39-47.
42. McCarthy AE, Lafleur C, Sutherland J, Lam P-P, Roth V, O'Connor AM, et al. Helping healthcare workers decide: evaluation of an influenza immunization decision tool. Can J Infect Control. 2010;25(1):21-4.
43. Smith SK, Trevena L, Simpson JM, Barratt A, Nutbeam D, McCaffery KJ. A decision aid to support informed choices about bowel cancer screening among adults with low education: randomised controlled trial. BMJ. 2010;341:c5370.
44. Légaré F, Ratté S, Stacey D, Kryworuchko J, Gravel K, Graham ID, et al. Interventions for improving the adoption of shared decision making by healthcare professionals. Cochrane Database Syst Rev. 2010(5):CD006732.
45. Vaughan SK. The importance of performance assessment in local government decisions to fund health and human services nonprofit organizations. J Health Hum Serv Adm. 2010;32(4):486-512.
46. Wunderlich T, Cooper G, Divine G, Flocke S, Oja-Tebbe N, Stange K, et al. Inconsistencies in patient perceptions and observer ratings of shared decision making: the case of colorectal cancer screening. Patient Educ Couns. 2010;80(3):358-63.

47. Mourits MCM, van Asseldonk MAPM, Huirne RBM. Multi Criteria Decision Making to evaluate control strategies of contagious animal diseases. Prev Vet Med. 2010;96(3-4):201-10.
48. Nobre FF, Trotta LT, Gomes LF. Multi-criteria decision making--an approach to setting priorities in health care. Stat Med. 1999;18(23):3345-54.
49. Ahmad F, Skinner HA, Stewart DE, Levinson W. Perspectives of family physicians on computer-assisted health-risk assessments. J Med Internet Res. 2010;12(2):e12.
50. Yancey AK, Cole BL, McCarthy WJ. A graphical, computer-based decision-support tool to help decision makers evaluate policy options relating to physical activity. Am J Prev Med. 2010;39(3):273-9.
51. Curtin JC, Lorenzi MV. Drug Discovery Approaches to Target Wnt Signaling in Cancer Stem Cells. Oncotarget. 2010;1(7):563-77.
52. Driscoll T, Dyer MD, Murali TM, Sobral BW. PIG--the pathogen interaction gateway. Nucleic Acids Res. 2009;37(Database issue):D647-50.
53. Eakin GS, Behringer RR. Diversity of germ layer and axis formation among mammals. Semin Cell Dev Biol. 2004;15(5):619-29.
54. Eisenberg JNS, Scott JC, Porco T. Integrating disease control strategies: balancing water sanitation and hygiene interventions to reduce diarrheal disease burden. Am J Public Health. 2007;97(5):846-52.
55. Jackson JA, Simon D, Greenbaum LA. False-positive sweat chloride test in a child with pyelonephritis in a single kidney. Clin Pediatr (Phila). 2009;48(6):683-5.
56. Lasher CD, Rajagopalan P, Murali TM. Discovering networks of perturbed biological processes in hepatocyte cultures. PLoS ONE. 2011;6(1):e15247.
57. Warren AS, Archuleta J, Feng W-C, Setubal JC. Missing genes in the annotation of prokaryotic genomes. BMC Bioinformatics. 2010;11:131.
58. Eisenberg JNS, Lewis BL, Porco TC, Hubbard AH, Colford JM. Bias due to secondary transmission in estimation of attributable risk from intervention trials. Epidemiology. 2003;14(4):442-50.

APPENDICES

Appendix A. Supplemental Information for Chapter I Manuscript

Modeling targeted layered containment of an influenza pandemic in the United States

Barrett *et al.* 10.1073/pnas.0706849105.

Online: http://www.pnas.org/content/105/12/4639/suppl/DC1

Supporting Information

Files in this Data Supplement:

SI Text
SI Figure 5
SI Table 5
SI Figure 6
SI Figure 7

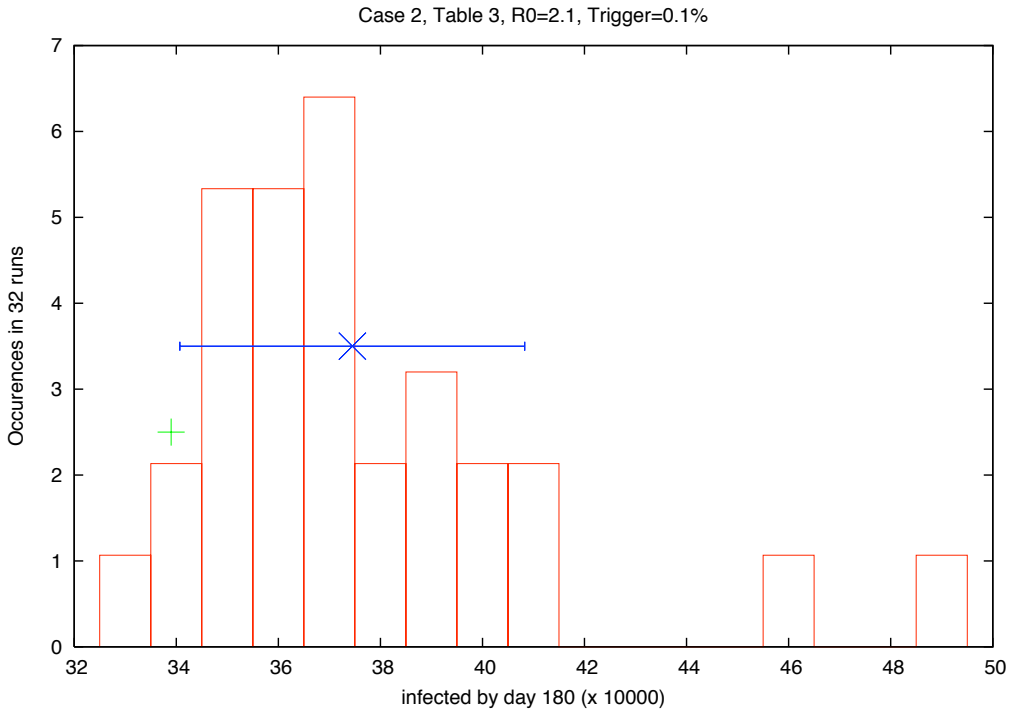

SI Figure 5

Fig. 5. Variability of VBI model over 32 runs, scenario 2, with a threshold of 1%, and R_0 = 2.1. The X denotes the mean, and the bar denotes one standard deviation on each side.

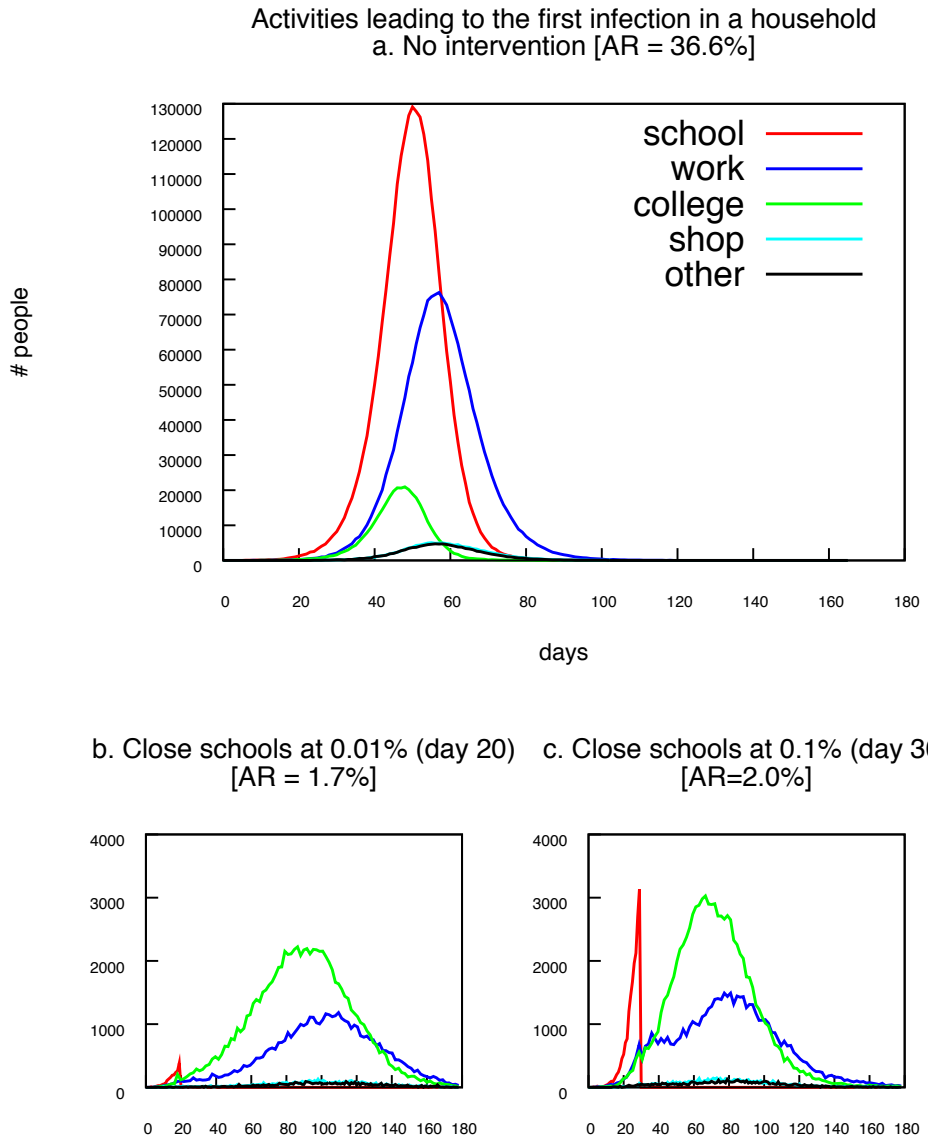

SI Figure 6

Fig. 6. Relative contributions of each activity type in the VBI model to interhousehold transmission, no intervention and scenario 2.

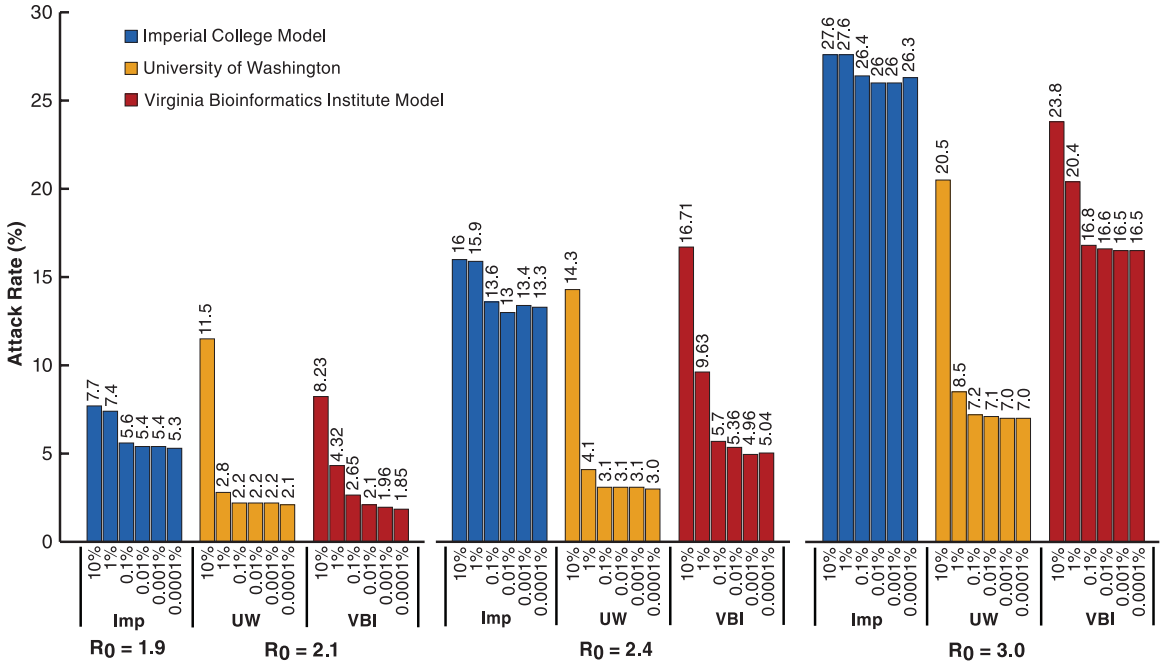

SI Figure 7

Fig. 7. Sensitivity to school closing threshold. Scenario 2 and three R_0 values, with only the school-closing threshold being varied between 10% and 0.0001% cumulative illness attack rates. Chicago population.

Table 5. UW/LANL model results

	R_0 =2.1	R_0 =2.4	R_0 =3.0
Chicago 10%	20.3 (0.57)	24.1 (1)	30.9 (0.49)
	201.1 (2.9)	262.9 (5.1)	371 (2.5)
Chicago 1% (base)	2.8 (0.11)	4.1 (0.2)	8.5 (0.74)
	38 (1.6)	61.2 (2.7)	138.4 (11.5)
Chicago 0.1%	0.32 (0.01)	0.44 (0.03)	1.3 (0.06)
	4.5 (0.14)	6.7 (0.41)	22.4 (1)

Chicago 0.01%	0.05 (0)	0.07 (0.01)	0.21 (0.03)
	0.77 (0.08)	1.2 (0.12)	3.7 (0.65)
Chicago 0.001%	0.02 (0)	0.02 (0)	0.06 (0.01)
	0.34 (0.03)	0.43 (0.02)	1.1 (0.21)

Mean of five realizations (standard deviation). Illness attack rates in first row; courses per 1,000 people in the second row.

SI Text

Description of the Three Models

The social structure of each of the three models was constructed differently, with implications for the effects of interventions. In the UW/LANL model, described in ref. 1, the population is geographically distributed among census tracts to closely represent variation in actual age and spatial distribution according to publicly available 2000 U.S. Census data (2). Each tract is in turn organized into 2,000-person communities as in refs. 3-5. The household size distribution corresponds to that in the 2000 U.S. Census data. There are eight mixing groups within which individuals can associate and be infected by others in these groups. Each person belongs to a household, neighborhood, a cluster of four households, and a community. School children go to primary school, middle school, or high school. Preschool children go to small day care centers or play in small play groups. Most adults go to work at workplaces. The neighborhoods and communities in this model provide the source of casual contact, say while shopping or going to a theater, whereas in the other contexts relatively close person-to-person association regularly occurs. Daytime contacts occur in neighborhoods and communities as well as in the age-appropriate setting, whereas nighttime contacts occur only in households, household clusters, neighborhoods, and communities. U.S. 2000 census data on tract-to-tract worker flow is used to configure the commute of working adults to their workplace, thus accurately capturing the short-to-medium distance population mobility important for disease spread. In addition, each individual takes occasional long-distance trips (3 per year on average), lasting between 1 day and 3 weeks (4.1 days on average), matching Bureau of Transportation Statistics data. School closure is implemented by closing the schools, day care centers, and small play groups, and doubling the household transmission probabilities. In the national models, schools are closed according to local illness attack rates, and the other thresholds are national illness attack rates. The model runs in cycles of two 12-hour periods (day and night). The results are reported at day 180 of the epidemic.

In the Imperial/Pitt model, described in ref. 6, individuals are colocated in households, with households being constructed to reflect typical generational structure

while matching empirical distributions of age structure and household size for the United States (7, 8). Households are randomly distributed in the modeled geographic region, with a local density determined by the Landscan population density dataset, which has subkilometer resolution (9). Exposure to infection comes from households, schools, workplaces, and random contacts in the community. The last of these depends on distance by using a gravity model, representing random contacts associated with movements and travel, and is the only means by which infection can cross national borders. The spatial kernel used to describe community transmission was parameterized from data on travel patterns in the United States. Air travel was also explicitly represented for national scale model runs. Data on the location, size, and grade composition of every school in the United States (10) were used in initializing the model, and children were allocated to schools with a gravity model so as to match empirical distributions of the journey distance between homes and schools. Similarly, data on commuting distances and the distribution of workplace sizes were used to configure the representation of workplaces in the model and to allocate individuals to workplaces. Thus, the community structure arises naturally from the variation in population density between urban and rural areas. The results at 180 days after the beginning of the epidemic are reported.

The VBI model has smaller mixing groups that are location-specific and determined by the types of activities people do as described in ref. 11, with further background on the concept of location-specific activities in ref. 12. The VBI model was developed by using U.S. Census data from Chicago tracts as well as the Public Use Microdata Samples (PUMS) (13) to build the synthetic Chicago data. Then activity data from the Chicago Area Transportation Study (CATS) (14) were used to synthesize regional travel patterns. The enrollment figures for each school were available from the Department of Education for public schools (10). Each person has a schedule of activities that are matched to activities reported by real people with similar demographics in activity surveys. For example, one person might go from home to work to a restaurant to night school and return home; another might go from home to school, several after-school activities, and home; and yet another might simply stay home all day. In this population, people have on average 4.66 activities per day, ranging from 1 to 22. The locations where these activities take place are chosen based on a form of a gravity model (11, 12) that statistically matches travel distances, by activity type, to those of real people from data gathered by regional transportation surveys. This process generates a realistic distribution of contact durations across age groups, activity types, and locations. For example, each person has from 0 to 262 (mean: 26) person-hours of contact at home and each worker has from 0 to 387 (mean 114) person-hours of contact at work. As each individual's schedule is manipulated to implement nonpharmaceutical interventions as described below, the distribution of contact durations changes, thus changing the population-wide mixing patterns. In this study, the probability of transmission between contacts is based solely on the resulting calibrated contact durations. Given the contact patterns, fitting the transmission rate-a single parameter-to desired values of R_0 yields age-specific attack rates between those of the 1957 and

1968 pandemics as well as the activity-specific attack rates shown in Table 3. The results at day 180 of the epidemic are reported.

In contrast to the other two models, the VBI model has colleges, shopping places, as well as other specific locations where the general community mixes. School closures and keeping children home are implemented on a per-household basis. In compliant households, an adult stays home with the child. In noncompliant households, the child's activity that immediately follows school was prolonged to last the length of school to simulate the mixing that an unsupervised child would do when schools are closed. The degree of interaction at colleges is not affected by social distancing measures.

In these models, R_0 is an output that is computed as part of a complex calibration that includes qualitatively matching to an age-specific illness attack rate pattern between those of the 1957 and 1968 pandemics.

Some Contrasting Implementations. Household transmission and model calibration: All three models base the assumptions about the home transmission on similar sources (15, 16), but used the information differently. The UW/LANL model uses the age-specific household secondary attack rates to calculate the daily probability of an effective contact within the household, then the other effective contact probabilities are calibrated to yield the desired age-specific attack rate pattern corresponding to that used by all three groups as the target pattern. A parameter corresponding to the transmissibility of the assumed pandemic strain is then varied, which yields the different target R_0 values. In the VBI model, the probability of transmission from an infectious person to a susceptible person during a contact (colocation) of duration that is given by the formula $ptrans = 1 - (1 - q)^t$, where q is a parameter that can be set independently in different locations, for different infectious or susceptible people. For each simulation run here, however, we have used the same value of q for every person and for every location. Differences between within-household and community infection rates in this model are due to different durations and different numbers of contacts in each venue. The overall R_0 is adjusted by changing the value of q. This can in principle affect the proportion of intrahousehold and community infection rates. VBI performed the calibration at a nominal $R_0 = 2.1$, but the difference in proportions as we change R_0 is minimal. In the Imperial/Pitt model, analysis of household infection data indicates that 30% of influenza transmission occurs in household. In the absence of data to inform the choice, transmission in other contexts was arbitrarily partitioned to give levels of within-place transmission comparable with household transmission, namely 33% of transmission was assumed to occur in schools and workplaces, and 37% in the wider community (i.e., in contexts other than households, schools, and workplaces). The transmission coefficient within schools was assumed to be twice that of workplaces. When varying R_0, the relative proportions of household, place, and community transmission were kept fixed.

Home isolation of cases: In the UW/LANL and VBI models, the compliance percent of cases are completely isolated, whereas the remaining percent continue circulation in the

community, as usual. In the Imperial/Pitt model, ascertained cases reduce their contact rates outside the home by the compliance percent.

School closing. In the VBI and Imperial models, an adult stays at home when the children in a household stay home. In noncompliant households in the VBI model, the child's activity that immediately follows school was prolonged to last the length of school to simulate the mixing that an unsupervised child would do when schools are closed. In the UW/LANL and Imperial/Pitt models, when children were not compliant, they continued to mix in the other mixing groups.

Quarantine of household contacts. Either 30 (60,90)% of the contacts are completely compliant with quarantine with the household contact rate doubling (UW/LANL) or contacts outside the home are reduced by 30 (60,90)% for the quarantined individuals (Imperial/Pitt). The VBI model changes the set of activities in which people participate, still providing activities for an entire day. The resulting change in the duration of contacts with other people at each activity determines the probability of transmission associated with that activity.

Workplace social distancing. In the UW/LANL and Imperial/Pitt models, workplace social distancing is implemented by reducing the contact probabilities. In the VBI model, workplace social distancing is achieved by reducing the maximum number of people allowed to be in the same sublocation, such as an office suite, from 50 to 25, representing a reduction in meeting sizes and other contacts.

Community social distancing. In the UW/LANL model, community social distancing is implemented by reducing the contact probabilities by a certain percentage in the community, the neighborhood, and the household clusters. In the Imperial/Pitt model, it is implemented by reducing the contact probabilities in the community. In the VBI model, community social distancing is implemented by removing the specified percentage of the appropriate activities (shopping, social and recreation, visits, and other).

Seeding of infectives. For the UW/LANL U.S. model, the 14 largest international airports act as gateways. Each day a random number (N) of infections is introduced into a random census tract near one of the 14 airports. N is proportional to the number of arriving daily passengers. For the Imperial/Pitt model the infection seeding is based on a global SEIR model that supplies external travelers to the United States at a rate that is based on an external epidemic that has an $R_0 = 1.6$. This model was used to calculate incidence of infection through time based on 73 million trips per year (Non-U.S. nationals plus U.S. travelers returning from overseas travel). Non-U.S. nationals select destinations proportional to population size. One third (1/3) of the trips are U.S. returnees. Also trips occur at locations according to population density. Returning travelers are expected to have more random destinations than foreign visitors. For the VBI Chicago model five people chosen uniformly at random from the entire population are used to seed the epidemic every day for the duration of the simulation. The same

sequence is used on every run. If one of the seeded persons has been prophylaxed or previously infected, no replacement is used.

SI Fig. 5 shows the relative contributions of each activity type in the VBI model to interhousehold transmission in the absence of intervention and in scenario 2. The first person to be infected in each household must acquire the infection outside the household at one of his/her daily activities. Each curve in SI Fig. 5 represents the number of first infections in a household acquired at the indicated activity. In an unmitigated epidemic, (SI Fig. 5*a*), schools stand out as the driver of interhousehold infection, responsible for 65% of all initial household infections by day 40. When this source is removed by closing schools, other mixing groups, such as colleges in SI Fig. 5*b*, emerge as the drivers, although at a much reduced level. SI Fig. 5*c* shows, however, that when schools stay open longer they can force transmission in the workplace, but workplace transmission does not appear to be a driver by itself.

Stochastic variability could explain some of the differences in the absolute illness attack rates under some scenarios among the models, because the number of realizations is relatively small. However, the population of 8.6 million in the Chicago population is relatively large, and the population is seeded regularly with infectives, so that the variability of the effectiveness measures of any given model is relatively small. SI Table 5 contains the mean and standard deviation of five realizations for the UW/LANL Chicago model for several of the scenarios. The variability is considerably under 5% for most scenarios, and nearly zero for some. Variability of the 10 realizations of the Imperial/Pitt model is not shown, but is similar to that of the UW/LANL model. The VBI did 32 runs of one scenario to examine stochastic uncertainty. The histogram based on 32 runs for scenario 2 is in SI Fig. 6. An examination of variability by VBI of 32 runs of scenario 2, $R_0 = 2.1$, with a 0.1% threshold yielded a mean of 3.75% (standard deviation of 0.35%) illness attack rate, with a range of one standard deviation on either side from 3.4% to 4.1%. With a baseline illness attack rate of 44.7%, the stochastic variability of the effectiveness is from 91.1% to 89.3%. The variability of the UW/LANL and Imperial/Pitt models is less. The VBI model requires much more computing time and resources, and they found it unfeasible to do multiple runs for each scenario and all sensitivity analyses.

Results

In a sensitivity analysis using the UW/LANL model, varying compliance in all three aspects of isolation, quarantine, and children staying home after school closure from 30% through 90%, with scenario 2 at a R_0 of 2.1, the illness attack was only reduced from 2.8% to 2.6%, indicating little sensitivity to the level of compliance in this scenario.

Discussion

Some of the interventions, such as home isolation of identified cases and home quarantine of their household contacts, depend on rapid ascertainment of symptomatic cases, whereas others, such as closure of schools and isolation of children in the home, reducing contacts in the workplace, and reductions in community contacts outside the home, depend on more general triggers.

The effectiveness of the interventions differs more among the three models at the higher R_0 of 3. The UW/LANL model is the most optimistic, likely because of the combination of the social structure and the slightly longer generation time compared with the Imperial/Pitt, so that both the general and the targeted interventions have a strong effect, the combination of which becomes more important at the higher R_0 values.

Models are generally constructed to explore specific scientific questions of interest. Different kinds of models than those presented here could be used to address other questions. Gani *et al.* (17) use a deterministic, age-structured model with groups at low and high risk of hospitalization to examine the optimal distribution of antivirals in reducing hospitalizations. They do not have households, workplaces, and schools and could not examine the scenarios we consider, but we do not have groups at high risk for hospitalization, so that is not a question that we explored. Colizza*et al.* (18) used a global model with travel between otherwise homogeneously mixing populations to explore the use of antivirals and travel restrictions to mitigate the pandemic. In contrast, our models assume that global spread has already occurred.

References:

1. Germann TC, Kadau K, Longini IM, Macken CA (2006) Mitigation strategies for pandemic influenza in the United States. *Proc Natl Acad Sci USA*103:5935-5940.
2. National Statistics Office (2000) Population and housing census 2000. Available at: http://www.nso.go.th.
3. Longini IM, *et al.* (2005) Containing pandemic influenza at the source. *Science* 309:1083-1087.
4. Halloran ME, Longini IM, Nizam A, Yang Y (2002) Containing bioterrorist smallpox. *Science* 298:1428-1432.
5. Longini IM, Halloran ME, Nizam A, Yang Y (2004) Containing pandemic influenza with antiviral agents. *Am J Epidemiol* 159:623-633.
6. Ferguson NM, Cummings DAT, Fraser C, Cajka JC, Cooley PC, Burke DS (2006) Strategies for mitigating an influenza pandemic. *Nature* 442:448-252, 10.1038/nature04795.
7. U.S. Bureau of the Census (2005) U.S. interim population projections . Available at: http://www.census.gov/ipc/www/usinterimproj/.
8. U.S. Bureau of the Census (2005) Households by size, 2005. Available at: http://www.census.gov/population/socdemo/hh-fam/tabHH-4.pdf.
9. Oakridge National Laboratory (2003) Landscan global population data. Available at: http://www.ornl.gov/sci/gist/landscan.
10. National Center for Educational Statistics (2004) Common core of data. Available at: http://nces.ed.gov/ccd/bat/index.asp.
11. Lewis B, *et al.* (2007) Simulated pandemic influenza outbreaks in Chicago (Virginia Bioinformatics Institute, Virginia Tech, Blacksburg, VA), Technical Report NDSSL-TR-07-004.
12. Eubank S, et al. (2004) Modelling disease outbreaks in realistic urban social networks. *Nature* 429:180-184.

13. Public Use Microdata Samples. Sample data, 2000. Available at: www.census.gov/main/www/pums.html.
14. Chicago Area Transportation Study. Transportation data, 2006. Available at: www.catsmpo.com.
15. Longini IM, Koopman JS, Haber M, Cotsonis GA (1988) Statistical inference for infectious diseases: Risk-specified household and community transmission parameters. *A*m J *Epidemiol* 128:845-859.
16. Cauchemez S, *et al*. (2004) A Bayesian MCMC approach to study transmission of influenza: application to household longitudinal data. *Stat Med*23:3469-3487.
17. Gani R, Hughes H, Fleming D, Griffin T, Medlock J, Leach S (2005) *Emerg Infect Dis* 11:1355-1362.
18. Colizza V, Barrat A, Barthelemy M, Valleron A-J, Vespignani A (2007) *PLoS Med* 4:0095-0110.

Appendix B. Antiviral Medkit Study

Following the study described in the manuscript in Chapter I, plans were formed and policies were enacted. However, additional questions arose about how to best deploy the pharmaceutical component of an influenza epidemic response. While several collaborators within MIDAS could not respond with full-scale simulations, VBI was able to respond in a short time frame. This capability was aided by both the use of synthetic individual level population as well as a highly optimized epidemic diffusion algorithm that benefited from the assumption that high resolution data could be aggregated into an efficient format.

Description of assumptions and key questions for the anlaysis:

VBI model assumptions

The following are the assumptions that will go into the setup of the MIDAS simulations to be conducted by the Virginia Tech group. These assumptions were modified from the original "Targeted Layered Containment (TLC)" study assumptions to better match the current assumptions used by pandemic planning policy makers and reflect many of the issues discussed in the past couple weeks.

Disease Model for Pandemic Influenza

Initial Infections: 5 randomly chosen people a day will be "externally infected" to seed the epidemic in the simulation population (as was done in the original TLC study) representing the importation of disease from outside the population.

Latent period: non-infectious, average 1.9 days

Days	frequency
1	30%
2	50%
3	20%

Infectious period: constant infectiousness, average 4.1 days

Days	frequency
3	30%
4	40%
5	20%
6	10%

Diagnosis: 67% of those infected have symptoms and 60% of symptomatic individuals are diagnosed 24 hours after the onset of symptoms

Anti-virals: Diagnosed individuals will be treated with a course of antivirals (Tamiflu)
Treatment – 5-day course that reduces the risk of transmitting infection by 15%
Prophylaxis – 10-day course that reduces the risk of becoming infected by 60%
Based on: Halloran et al. *American Journal of Epidemiology*. 2007;165; 212-221

Changes to Human Behavior

Triggering event: A cumulative total of some percent of the population has been diagnosed with pandemic flu. The following interventions will be implemented once this triggering event has been reached representing the recognition the arrival of the pandemic and the initiation of pandemic response plans.

Non-Pharmaceutical Interventions (NPIs): These community mitigation strategies will take place following the triggering event defined above. These interventions will have limited compliance and are listed in order of priority such that not all interventions may be implemented (or lower priority interventions may have lower levels of compliance) to reach the "50% reduction in attack rate through community mitigation strategies" as has been requested.

Isolation – Compliant individuals diagnosed with influenza will be isolated from contact with all individuals other than their household members
Quarantine of household members – Members of households with a ill individual complying with isolation will also self-isolate and will only interact with other household members.
School closure – Interactions occurring through school activities will cease, compliant students will not interact with other students, while non-compliant students will have interactions with other students outside of school.
Liberal Leave – After 24 hours of symptoms, ill individuals will cease all work activities.
Community Social Distancing – Compliant individuals curtail non-essential activities (essential interactions occur at home, work, and school)
Work place social distancing – The likelihood of infections occurring in the workplace will be reduced by some percent.

Antivirals use by individual

Public stockpile: A fixed number of courses, corresponding to roughly 20% of the population (60-68 million US-wide distributed proportionally by population) will be made available for treatment.

Private stockpile: Selected individuals and their households will have access to pre-purchased antiviral medkits (totaling 5%, 15%, 25% of the entire population). How

these medkits are distributed throughout the population is the main experimental variable for this study, we intend to study 3-5 different schemes for all 3 different levels.

Use by individuals: Individuals diagnosed with disease will take a course of antivirals from their private stockpile if their household has pre-purchased the antiviral medkits and they haven't already used it, this will trigger the household members to use their antivirals prophylactically. If an individual doesn't have a private stockpile, they will use a course from the public stockpile if it hasn't been exhausted.

Misdiagnosed cases: Assuming other respiratory infections and a panicky population many people without influenza infection will be diagnosed (medically-assisted or by themselves) with pandemic influenza. How this is modeled is still a question (see question 2 below), however it is done we will assume that these misdiagnosed individuals will get their antivirals from the two sources in the same way as the correctly diagnosed individuals (described above).

General assumptions about model behavior

Population: Chicago population (8.8 million) used in previous study

Final attack rate of epidemic: 30% of population will have symptomatic infection (under current assumptions, 67% symptomatic, this means 45% of the population will become infected)

Calibration Process:

1. Adjust global transmissibility so that completely unmitigated epidemics have a clinical attack rate of 30% (45% infection rate with 2/3 symptomatic)
2. Layer interventions from the NPI list above one by one and adjust compliances (initially starting with levels used in Scenario 2 of the TLC study) until the 30% is reduced to 15%

Open Questions

1. Antivirals
    a. Efficacies of antivirals – What how effective should we consider these antivirals for both treatment and as a prophylaxis? There is variability in the reported efficacies, and even the meta-analysis has limited power, should we err on the conservative side?
    b. Triggered use – Is all antiviral use dependent on the trigger? Even private?
2. Misdiagnosis of influenza
    a. Impact on triggers – Should the misdiagnosed ILIs and other respiratory infections count towards the trigger?
    b. Treatment – If there is public stockpile available, do misdiagnosed cases use it as well? Does this only happen after the trigger?

c. Who can be misdiagnosed? Should those who have already been infected with the flu be able to get an ILI and be misdiagnosed? If you've been misdiagnosed previously and were treated from the public stockpile but then get infected with pandemic flu can you at get more antiviral?
d. How?
    i. Should this be a fixed number of people per day that are likely to have ILI and be diagnosed with the flu, set such that overall PPV averages out to 35%?
    ii. Should we pick a fixed proportion (1.86 for 35% PPV) of individuals to be misdiagnosed everytime a true influenza infection is diagnosed?
e. 35% of what exactly? Depending on the questions above the way that this ratio is calculated (and thus the numbers of misdiagnosed) could vary a bit, it would be helpful to clearly state what the ratio represents.

Results

Different scenarios simulated:

Scenario	Final Attack Rate	Description
Base	34%	Completely unmitigated epidemic
NPIs	22%	Non-Pharmaceutical Interventions are implemented, social distancing, school closure, etc.
Public	18%	MedKits are made available through a public stockpile (enough for 20% of the population) for treatment, first come first serve based on diagnosis
PPV	20%	Public stockpile with the effects of false positive diagnoses wasting
Private	19%	Private stockpiles, individuals pre-purchased medkits and use rapidly at home upon a household member experiencing symptoms

Table 1. Description of the different policies summarized in Figures 1 and 2

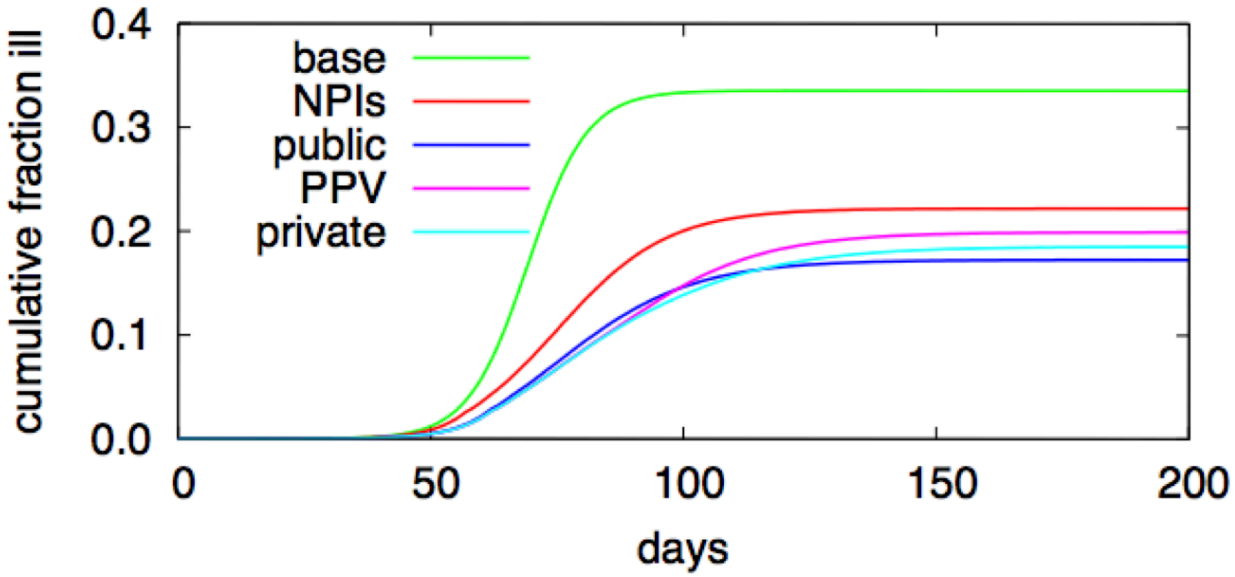

Figure 1. Cumulative numbers of infections under the different scenarios.

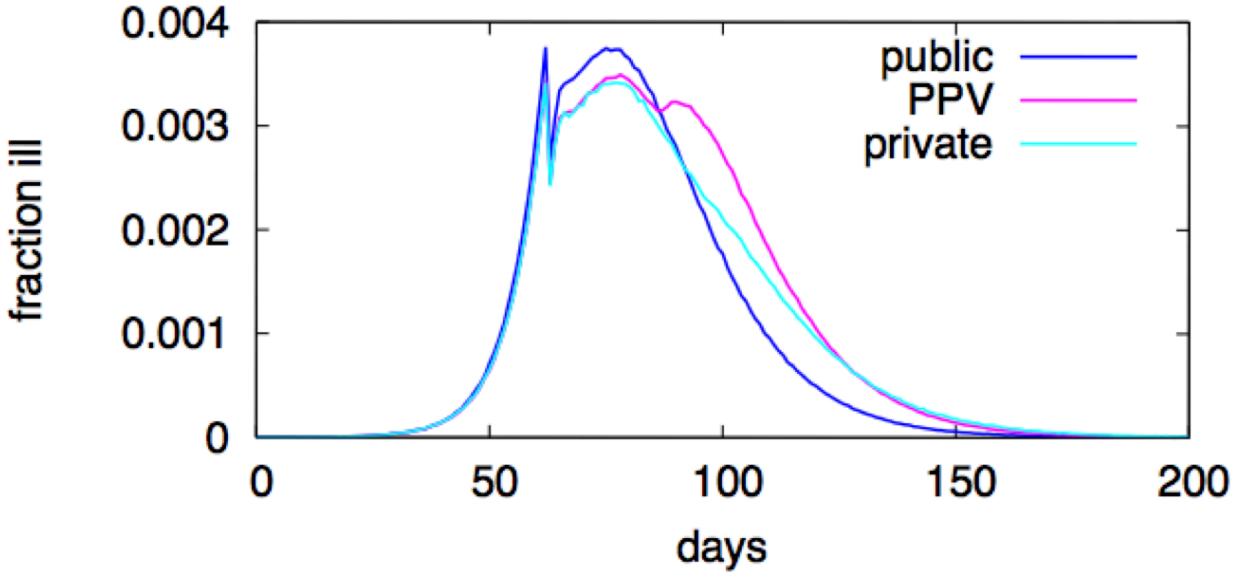

Figure 2. Daily fraction of population infected showing the effects of public and private stockpile use

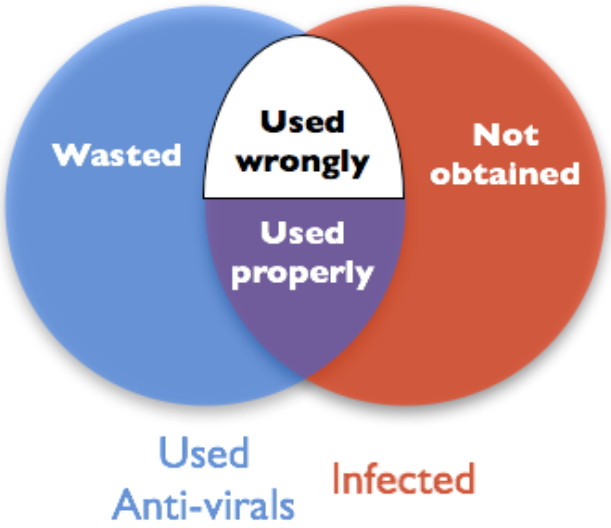

Figure 3. Venn Diagram illustrating how private stockpiles might be used. Not all individuals infected will have had the chance to purchase a stockpile, of those that have a stockpile and use the antivirals, some will be used wrongly, only a fraction will use properly

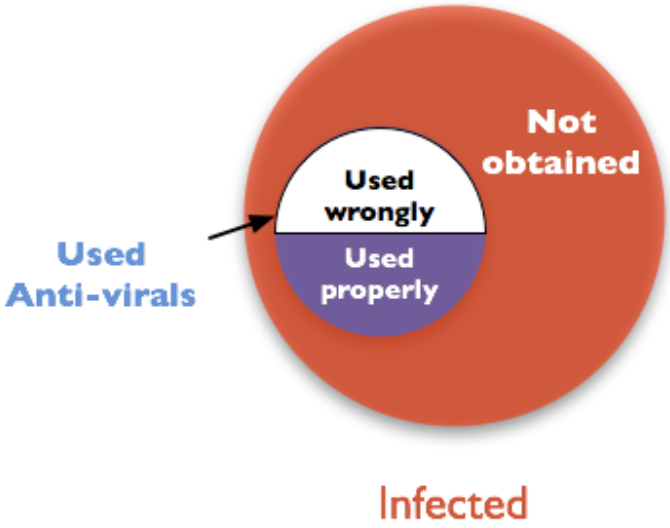

Figure 4. Venn Diagram illustrating how public stockpiles might be used. Since these are only adminstered to those with symtoms, most of the used anti-virals are amongst the infected

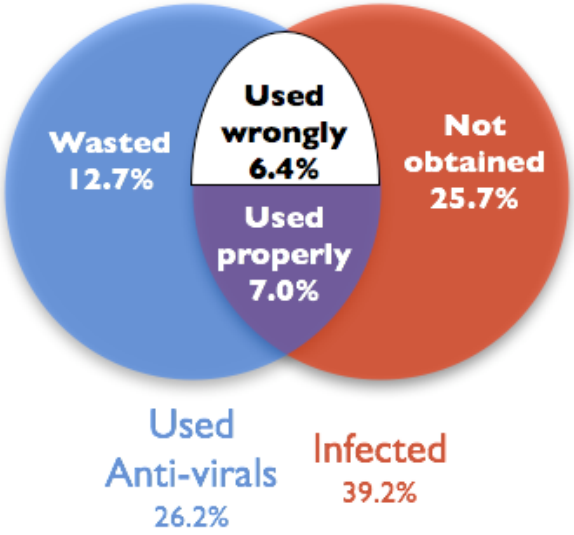

Figure 5. Quantification of the different groups based on simulations. Base on no private stockpiles

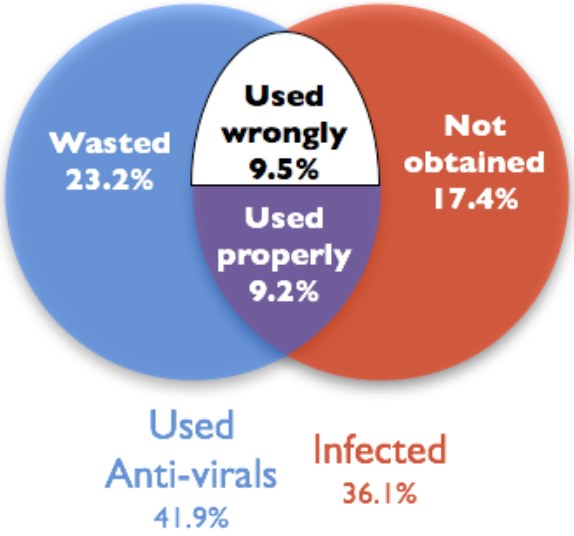

Figure 6. Quantification of the different groups based on simulations. Based on 25% of the available anti-virals being in private stockpiles distributed at random in the population.

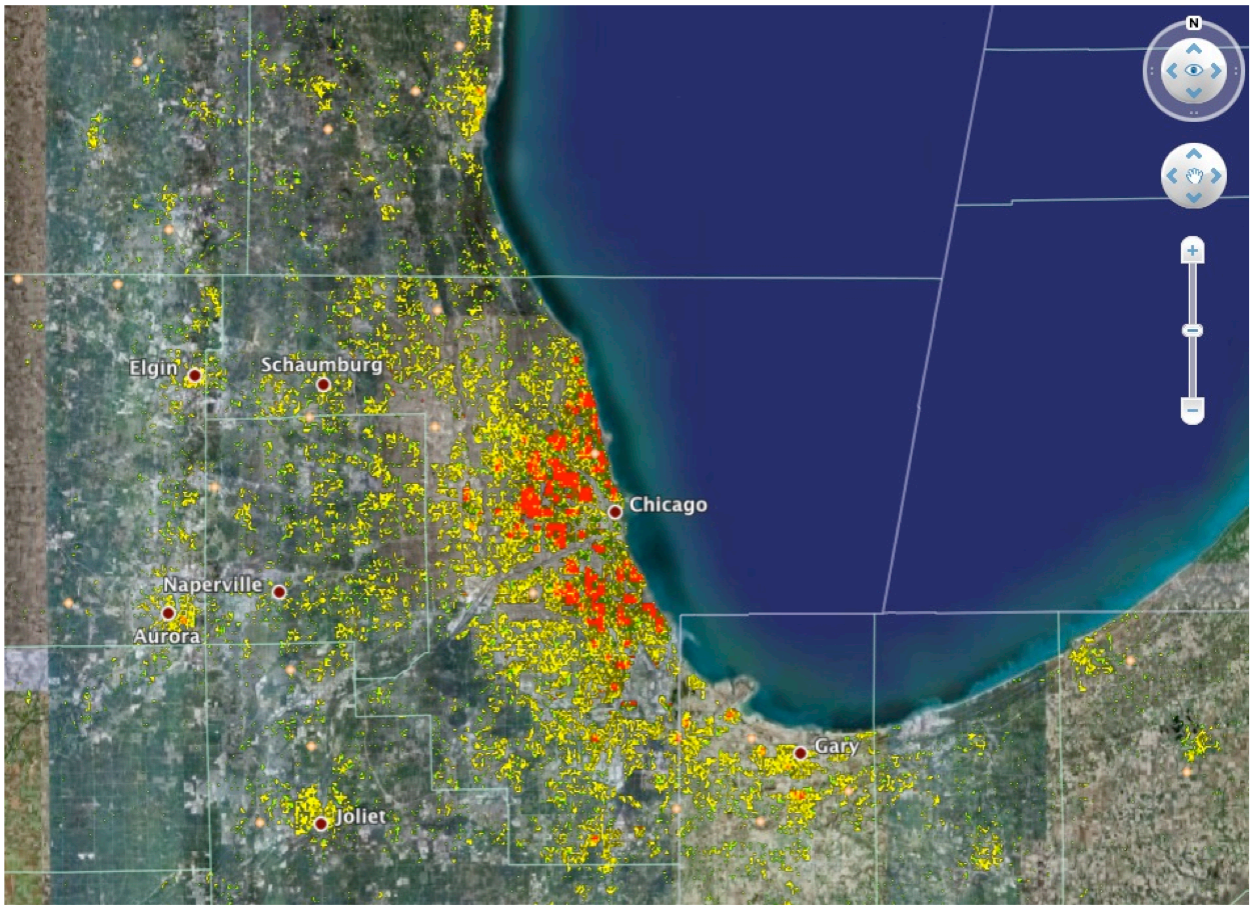

Figure 7. Geographic distribution of private stockpile when it is distributed to the low income individuals (as opposed to the proposed market-based distribution that would favor high income households).

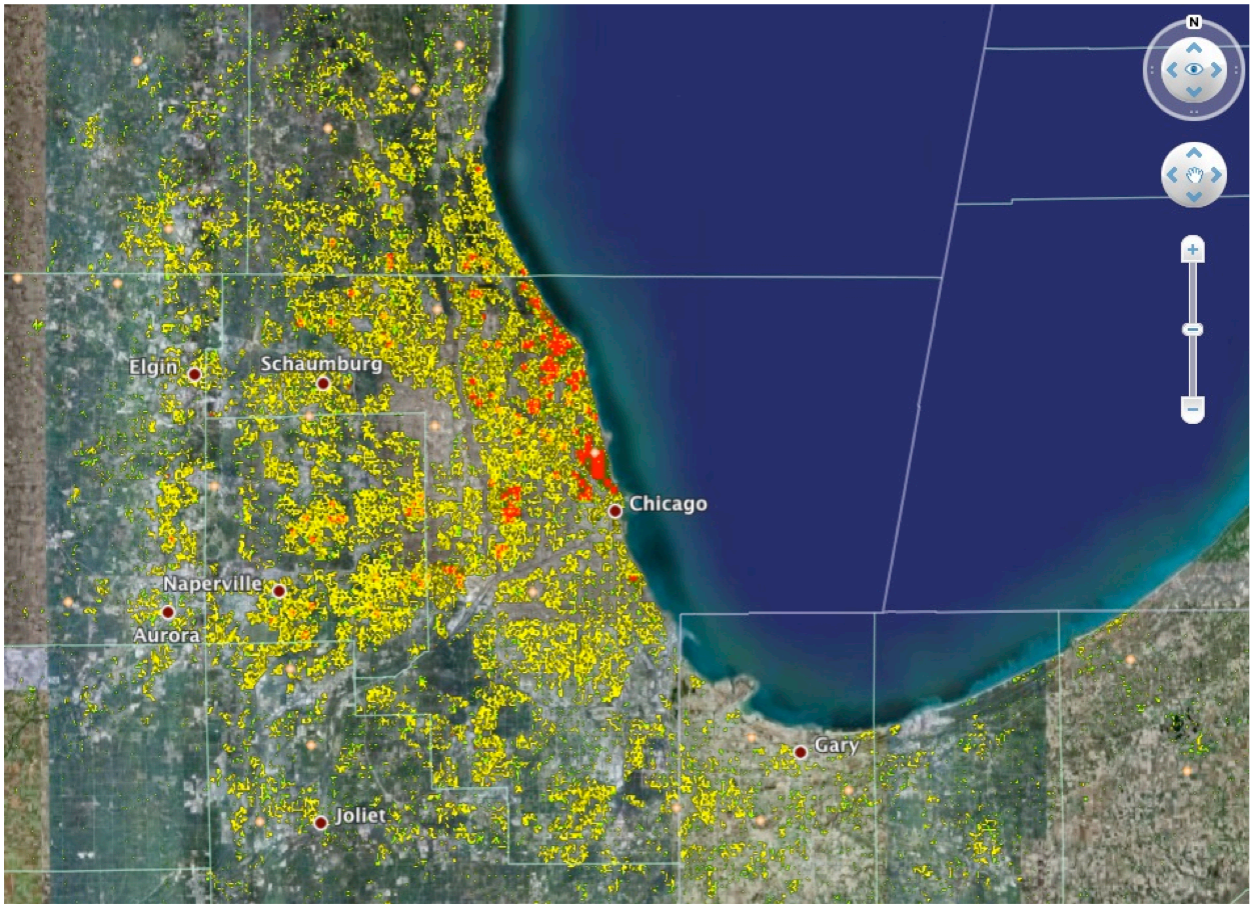

Figure 8. Geographic distribution of private stockpile when it is distributed to the higher income individuals (as is likely under the proposed market-based distribution).

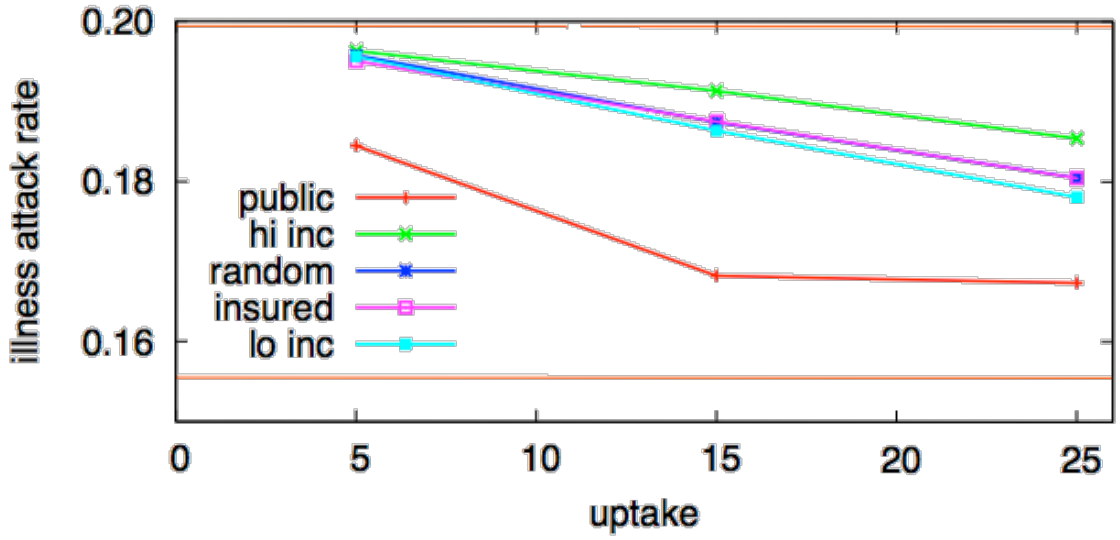

Figure 9. Attack rates based on uptake of the stockpile and its distribution. Shows that while there is only small overall differences (~2%) in attack rates, who gets the private stockpiles matters, foregoing any implications on the social justice of the decision. Public stockpiling is also shown to out-perform any private stockpile distributions.

Appendix C. H1N1 Emergence Study

Suspicions deepened that the an emerging strain of H1N1 influenza may cause widespread infections. Intense interest gathered to evaluate how severed and widespread the epidemic might be. Initial data surrounding the emerging disease were incomplete and efforts were focused on analyzing this data and estimating the extent of an outbreak based on a range of possible disease characteristics. Government analysts were able to use a web-based environment created by the VBI group to directly input these disease characteristics as they were revealed and keep high-level decision makers informed about possible scenarios within a 24 hour decision cycle. The software based on highly detailed synthetic populations allowed the deep involvement of the government analysts which accomplished response times not previously possible between government and academic institutions.

The following document is a rapid analysis conducted to inform government analysts the potential range of plausible disease models that fit the current estimates of H1N1 disease. It goes on to offer potential outcomes of the epidemic based on these models.

Disease models for H1N1, S-OIV

The novel H1N1 swine-origin influenza virus recently identified presents many challenges. The lack of any clear evidence surrounding its transmissibility, natural history, and severity make it particularly challenging to predict its potential impact with any confidence. This in turn makes decision making about the interventions to put in place challenging, since the benefit side of the cost/benefit ratio is so ill-defined.

Current computational models of this virus rely on "best guesses" as the basis for their calculations. The model underlying the simulations used in the DIDACTIC modeling platform created by the NDSSL at VBI, is based on the natural history of seasonal flu which is a reasonable place to start for an unknown strain of pandemic influenza. Now that a little more information is known a slightly more refined models can be developed. Once details of the contact investigations currently underway are better known, these models can be even more refined.

The key parameters in terms of how an epidemic spreads are incubation period, duration of infectiousness, and the transmissibility. Two different disease models were constructed following different estimates for their natural history. One was based on a description of a swine flu outbreak on an Army base in 1976 (Lessler, 2007). The other was based on a natural history that is implied by the conservative estimates of case contact identification provided by the CDC for the current H1N1 S-OIV pandemic (CDC, 2009). Both of these descriptions lacked information on the likeliness of an individual to experience the incubation and infectiousness durations specified, so the well-described distributions of the seasonal flu were used to interpolate the distributions.

CDC Guidelines

Incubation		**Symptomatic**	
Days	Proportion	Days	Proportion
1	0.1	3	0.2
2	0.15	4	0.3
3	0.2	5	0.2
4	0.2	6	0.15
5	0.2	7	0.1
6	0.075	8	0.05
7	0.075		
Expected average	**3.775**		**4.8**

Fort Dix

Incubation		**Symptomatic**	
Days	Proportion	Days	Proportion
0	0.4	2	0.4
1	0.6	3	0.5
		4	0.1
Expected average	**0.6**		**2.7**

Simulations were run using these two disease models with 3 different levels of transmissibility to explore the sensitivities to both the different natural histories and their relative strengths. The following results show that not only does the level of transmissibility play a role in the shape of the epicurve and the final attack rate, but the distribution of the incubation period and the infectious period. These effects are

significant as they can change the day of the peak number of people infected by over a hundred days. This highlights the importance of knowing as much as possible about the disease.

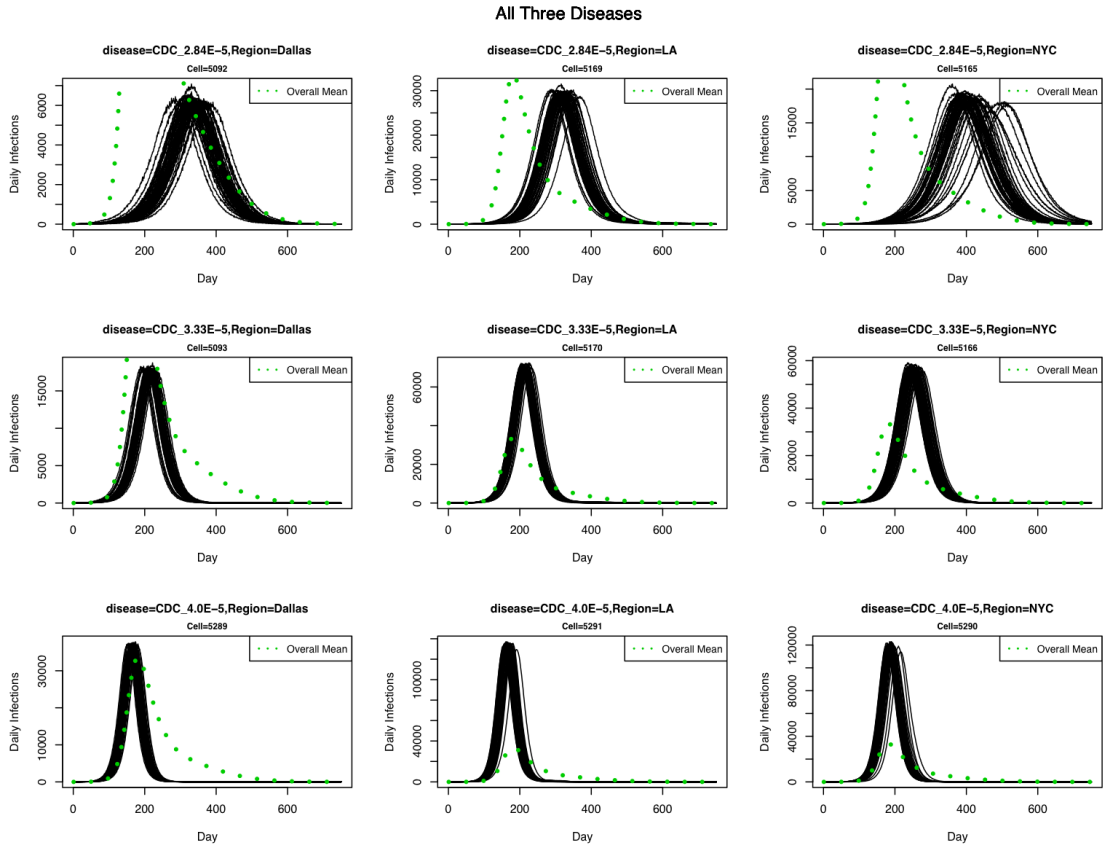

Figure 1. Epicurves for the CDC guideline based disease model, with 3 different levels of the transmissibility, in three different regions (Dallas, LA, NYC).

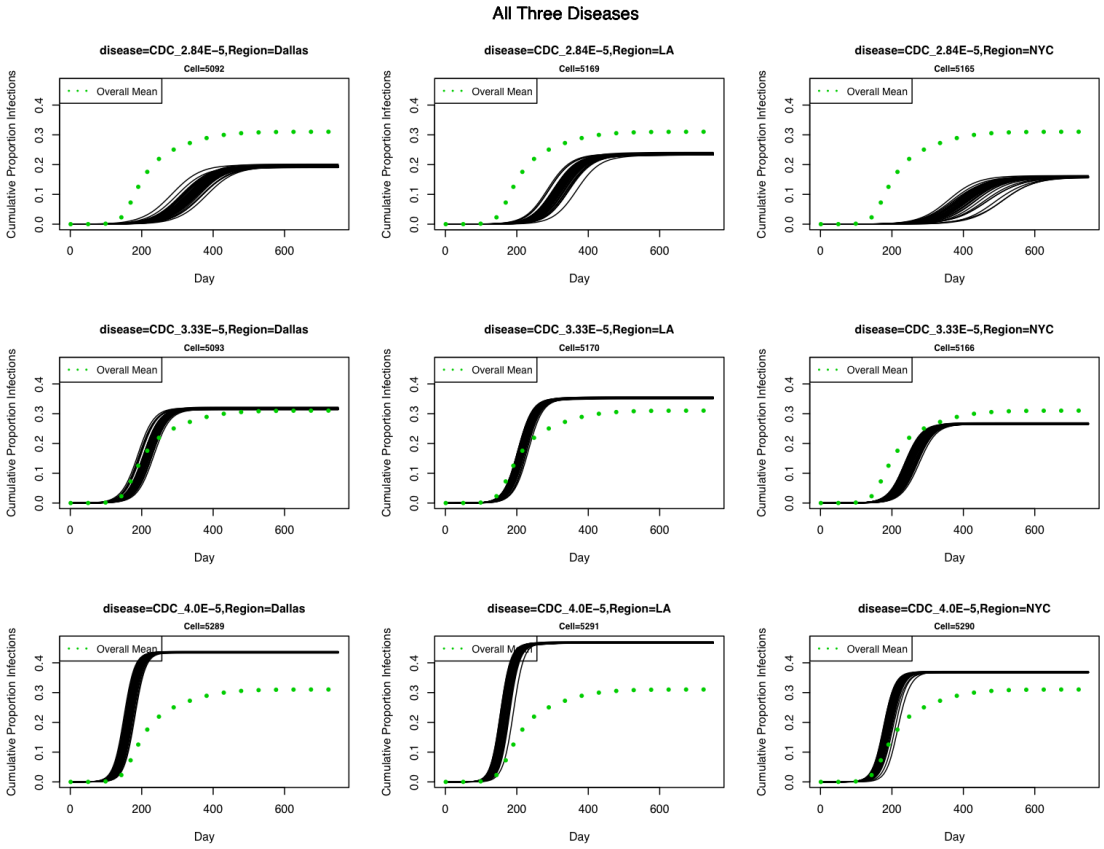

Figure 2. Cumulative Epicurves for the CDC guideline based disease model, with 3 different levels of the transmissibility, in three different regions (Dallas, LA, NYC). Shows the final attack size:

	Final Attack Rate		
Transmissibility	Dallas	LA	NYC
2.8e-5	19.34%	23.47%	15.85%
3.3e-5	31.70%	35.25%	26.61%
4.0e-5	43.52%	46.72%	36.82%

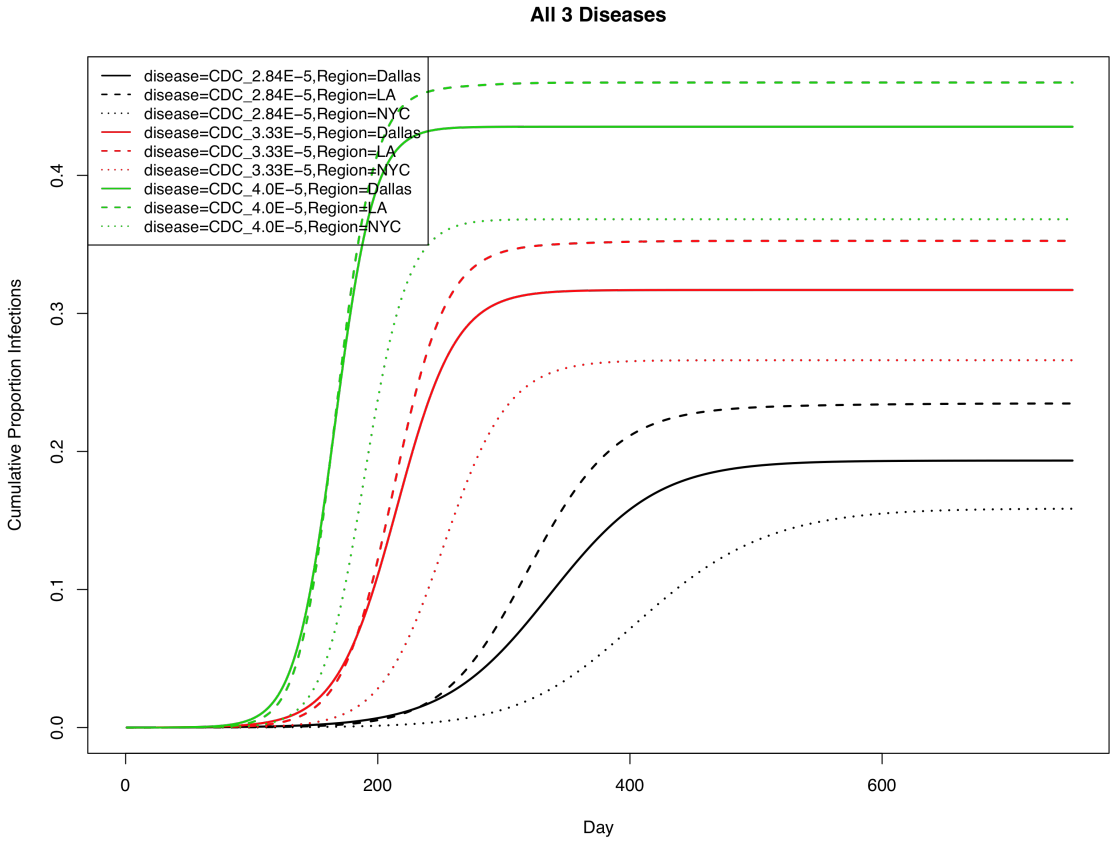

Figure 3. Cumulative Epicurves for the CDC guideline based disease model, with 3 different levels of the transmissibility, in three different regions (Dallas, LA, NYC), one plot.

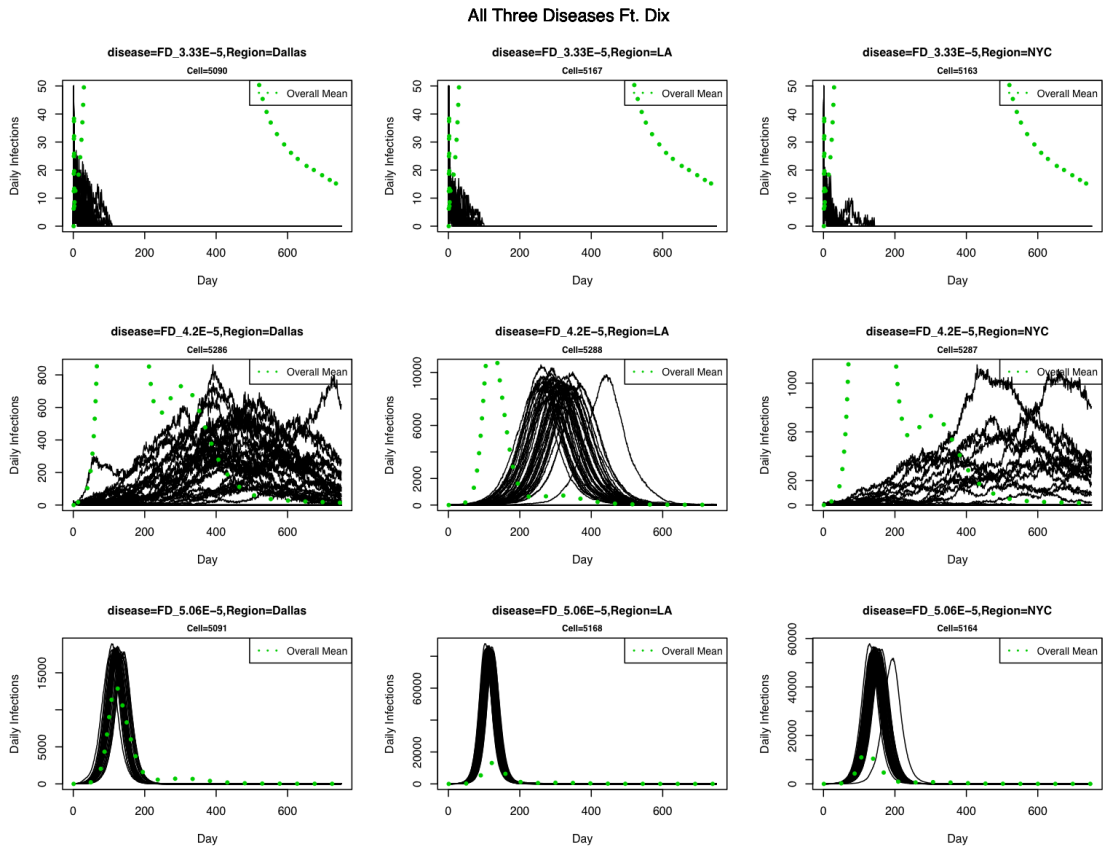

Figure 4. Epicurves for the Fort Dix study based disease model, with 3 different levels of the transmissibility, in three different regions (Dallas, LA, NYC).

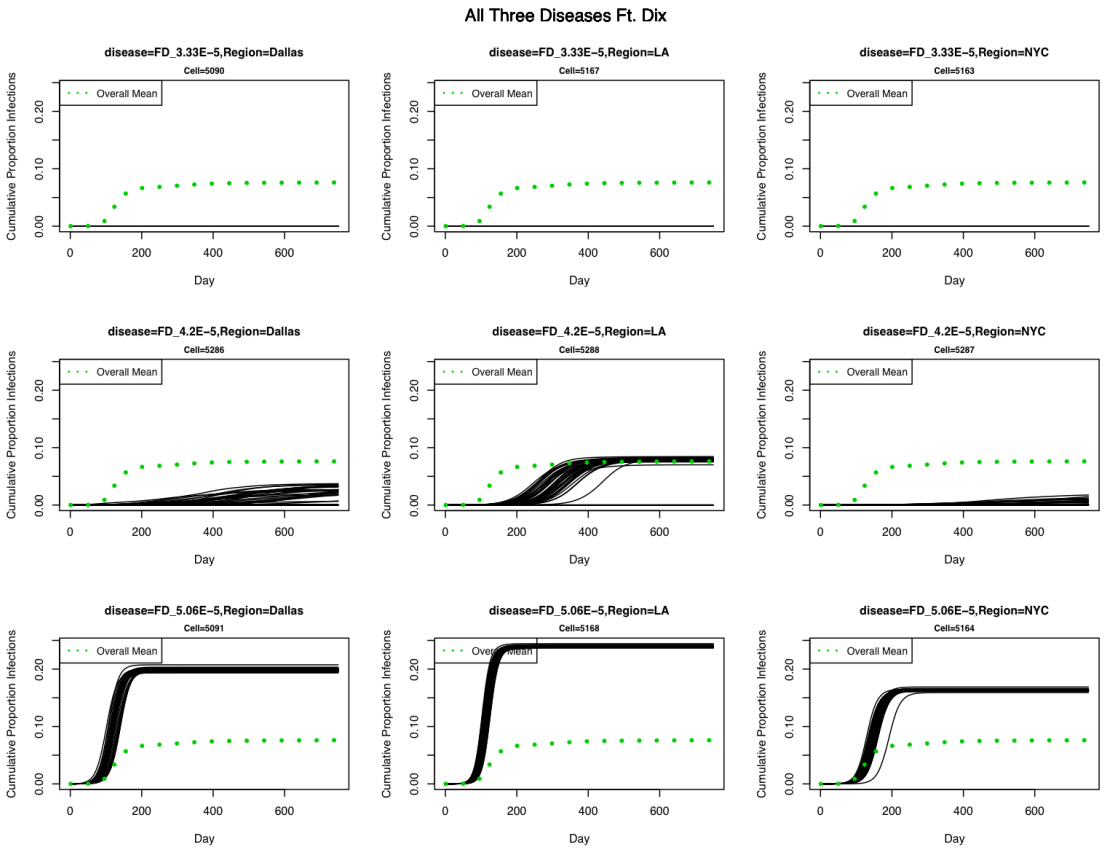

Figure 5. Cumulative epicurves for the Fort Dix study based disease model, with 3 different levels of the transmissibility, in three different regions (Dallas, LA, NYC).

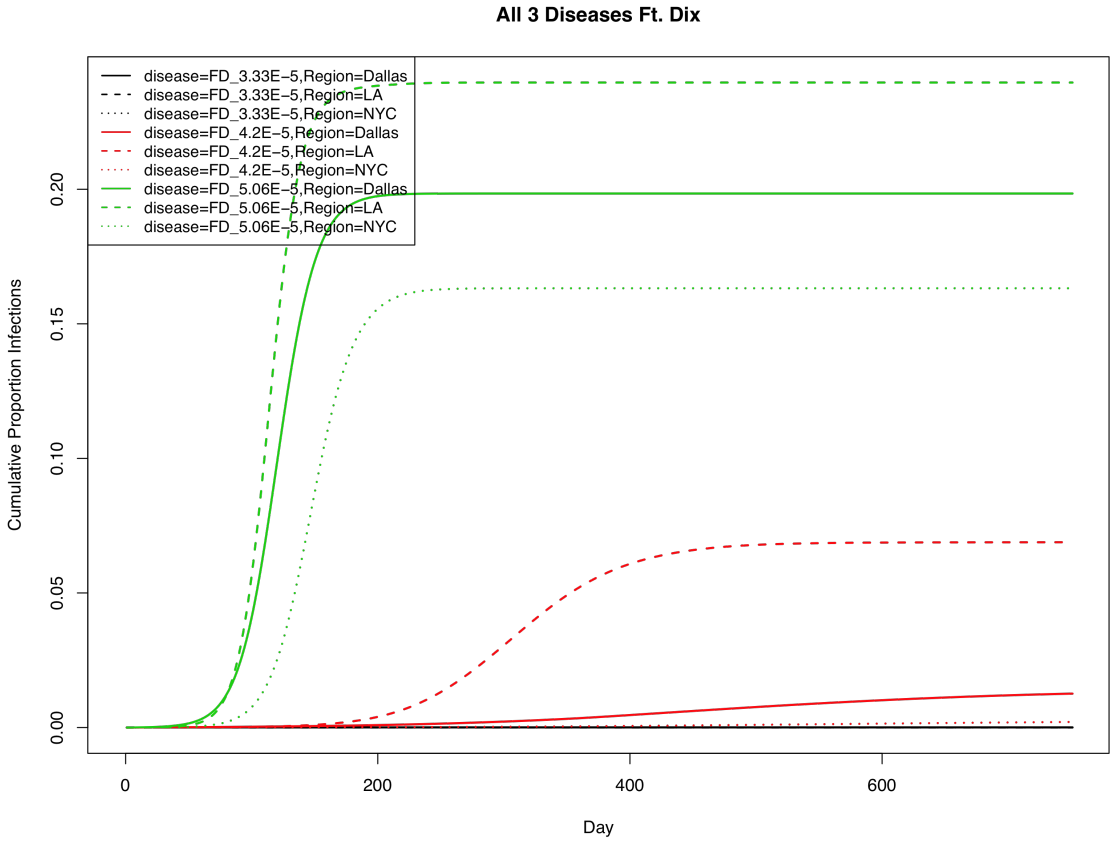

Figure 6. Cumulative epicurves for the Fort Dix study based disease model, with 3 different levels of the transmissibility, in three different regions (Dallas, LA, NYC), on one plot.

References

Lessler et al. Transmissibility of swine flu at Fort Dix, 1976. Journal of the Royal Society Interface (2007) 4, 755-762.

CDC. Interim Guidance on Case Definitions to be Used For Investigations of Swine-Origin Influenza A (H1N1) Cases. CDC website (April 27th, 2009) http://cdc.gov/h1n1flu/casedef_swineflu.htm

Appendix D. H1N1 Third Wave Study

Soon it became clear that the H1N1 influenza virus was going to cause a pandemic, and many of the policies already in place thanks to previous the planning efforts, were enacted. Nonetheless questions lingered and support from the modeling community was sought. In particular as it became clear that the epidemic had reached its natural peak before the bulk of the vaccines were able to be deployed and effective, concerns arose that reduced uptake of vaccine might leave a sufficient portion of the population susceptible and there was a risk of a "third wave" of infections following the Christmas and New Year holiday as had been previously experienced in other pandemics. In a week period a multi-factorial (48 cell) study was designed, implemented, analyzed, and relayed to decision makers.

The study design follows:

BARDA - 3rd wave modeling

Question: What conditions are necessary to permit a 3rd wave of H1N1 to occur?

Context: Multiple different conditions could conspire to create a 3rd wave of pH1N1. Previous simulations suggest that seasonal forcing alone is not sufficient to restart an epidemic in January given the current immunity conferred through vaccination, natural infection, and previous H1N1 experience (in the 1950s) remains perfect. It is impossible to anticipate whether the virus will drift sufficiently to make this immunity less than perfect. We will explore the resilience of the population to seasonal forcing, earlier onset of seasonal forcing, and loss of perfect immunity.

Simulation Parameters

Populations: "The Washingtons" ie DC and Washington state (WA pop 3.1)

Immunity:
1950's - Selection of random individuals above 65, 35% have immunity applied
1st wave - Selection of random individuals without regard to age, 0.3% have immunity
Vaccination - Selection of random individuals, in proportion to a weekly prorated level consistent with the vaccine delivery schedule. The effectiveness of the vaccine is delayed by 20 days (assuming 10 days from vaccine delivery to "in arms", and an additional 10 days for vaccine to produce immunity in the person).
2nd wave - Derived from simulations run with vaccine scheduled and then represented by taking the top proportion equal to the individuals that were most likely to be infected in the 2nd wave of infection.
The immunity from these factors are all assumed to be perfect or 50%. Partial immunity is modeled by allowing the probability of infection from the disease to be reduced by the

specified fraction.

Seasonality: Transmission probabilities becomes more likely as individuals susceptibilities are increased by 0%, 20%, and 40%, Globally and the onset of the seasonal change occurs on dates Dec 1st, Dec. 15th, and Jan 1st

Parameter Summary:
2 populations x 3 levels of seasonal forcing x 3 dates of seasonal onset x 2 levels of susceptibility = 36 cells

Baseline Simulation Description:
The H1N1 disease modeled is similar to seasonal flu (incubation period is assumed to have a 50% clinical attack rate and to cause an epidemic that infects 30%

Implementation:
Construct immunity profile for 1950s and first wave, assign immunity level by vaccination at time 0%
Construct vaccine delivery Triggers and Actions, properly delayed from rollout. Oct 18th is the beginning of the 43rd week which represents the peak of reported ILI, match peak of epidic to this (for trans 4.5E-5 ~130), back calculate to Oct 3rd when first vax was delivered, and then delayed, thus first vaccine efficacy is on day 136 (130 - 14 + 20).
Simulate 2nd wave with these characteristics.
Analyze 2nd wave simulation results to determine pool of individuals most likely to be infected in a 2nd wave.

Results

A presentation made to officials in the Biomedical Advanced Research and Development Authority in the Department of Health and Human Services on December 16th, 2009 in Washington DC.

Models of Infectious Disease Agent Study

MIDAS

Conditions for a 3rd Wave of pH1N1

Bryan Lewis
and
Stephen Eubank

Washington DC, Dec 16th, 2009

Simulation Parameters

- **Population**: Washington DC metro area
- **Disease**: Symptomatic and asymptomatic (50%), incubation 1-3 days, infectious 3-6 days
- **Treatment**: Diagnosed symptomatic infections treated with oseltamivir (70% compliance)
- **Software**: EpiFast initially configured using web-based DIDACTIC tool
- **Design**: Fully Balanced considering all combinations of 5 different parameters

VBI Virginia Bioinformatics Institute at Virginia Tech
Network Dynamics and Simulation Science Laboratory
Virginia Tech

Outline of Experimental Protocol

- Simulate 2nd wave
- Calculate immunity profile of population
  - Present: Dec 15th
  - Next year: Jan 15th
- Initialize simulation with immunity profile and a variety of different conditions
- Observe the size of "3rd wave"

2nd Wave Details - Vaccinations

- Vaccine administered on schedule of provided by BARDA
  - Adjusted to match delays in deliveries
  - Total delivered US-wide by Jan 2 – 60M or 30M
- Administered according to age
- Efficacy of 70%
- Delay of 20 days from allocation to effectiveness

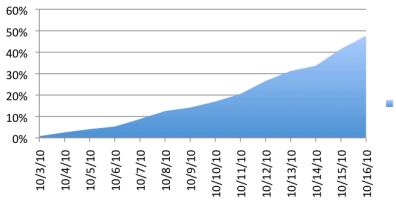

Age Group	Proportion of vaccinations
0 - 17	45%
18 - 64	46.7%
65 +	13.3%

VIRGINIA BIOINFORMATICS INSTITUTE AT VIRGINIA TECH
Network Dynamics and Simulation Science Laboratory
Virginia Tech

2nd Wave Details – Final Attack Rate

- Two Scenarios
  - Initial estimates of 2nd wave were ~30% (R0 of ~1.3-1.5)
  - CDC published estimates that put the national cumulative infections up to Nov 15th at 15%
    - To match this and a peak at Oct 14th based on ILI data the overall attack rate is smaller: ~21%

Simulated 2nd Waves aligned to ILI data

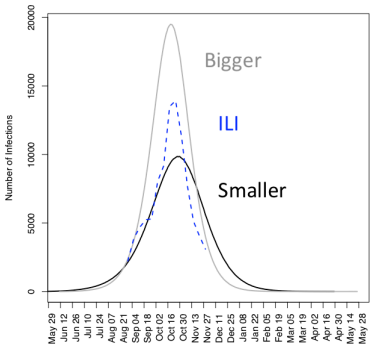

2nd wave Simulations

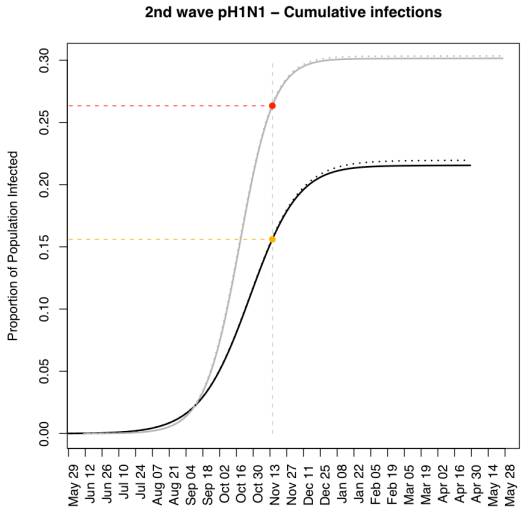

"Bigger"
Final infections = 30%
Nov 15th = 26%

"Smaller"
Final infections = 21.5%
Nov 15th = 15.6%

3rd Wave Conditions

- Loss of immunity from vaccination & previous infection
  - No change or 50% loss
- Size of 2nd wave
- Seasonal changes increase transmissibility
  - Increases of 0%, 20%, or 40%
- Timing of Changes
  - Dec 15th or Jan 15th
- Declining vaccine acceptance
  - "In arms" US wide 60M or 30M
- 48 individual scenarios (2x2x3x2x2) with 25 iterations each

Immunity Loss	2nd Wave Size	Seasonality	Date	Vaccine Acceptance
0%	30%	0%	Dec 15th	30M
0%	30%	0%	Dec 15th	60M
0%	30%	0%	Jan 15th	30M
0%	30%	0%	Jan 15th	60M
0%	30%	20%	Dec 15th	30M
0%	30%	20%	Dec 15th	60M
0%	30%	20%	Jan 15th	30M
0%	30%	20%	Jan 15th	60M
0%	30%	40%	Dec 15th	30M
0%	30%	40%	Dec 15th	60M
0%	30%	40%	Jan 15th	30M
0%	30%	40%	Jan 15th	60M
0%	21%	0%	Dec 15th	30M
0%	21%	0%	Dec 15th	60M
0%	21%	0%	Jan 15th	30M
0%	21%	0%	Jan 15th	60M
0%	21%	20%	Dec 15th	30M
0%	21%	20%	Dec 15th	60M
0%	21%	20%	Jan 15th	30M
0%	21%	20%	Jan 15th	60M
0%	21%	40%	Dec 15th	30M
0%	21%	40%	Dec 15th	60M
0%	21%	40%	Jan 15th	30M
0%	21%	40%	Jan 15th	60M
50%	30%	0%	Dec 15th	30M
50%	30%	0%	Dec 15th	60M
50%	30%	0%	Jan 15th	30M
50%	30%	0%	Jan 15th	60M
50%	30%	20%	Dec 15th	30M
50%	30%	20%	Dec 15th	60M
50%	30%	20%	Jan 15th	30M
50%	30%	20%	Jan 15th	60M
50%	30%	40%	Dec 15th	30M
50%	30%	40%	Dec 15th	60M
50%	30%	40%	Jan 15th	30M
50%	30%	40%	Jan 15th	60M
50%	21%	0%	Dec 15th	30M
50%	21%	0%	Dec 15th	60M
50%	21%	0%	Jan 15th	30M
50%	21%	0%	Jan 15th	60M
50%	21%	20%	Dec 15th	30M
50%	21%	20%	Dec 15th	60M
50%	21%	20%	Jan 15th	30M
50%	21%	20%	Jan 15th	60M
50%	21%	40%	Dec 15th	30M
50%	21%	40%	Dec 15th	60M
50%	21%	40%	Jan 15th	30M
50%	21%	40%	Jan 15th	60M

3rd Wave Results

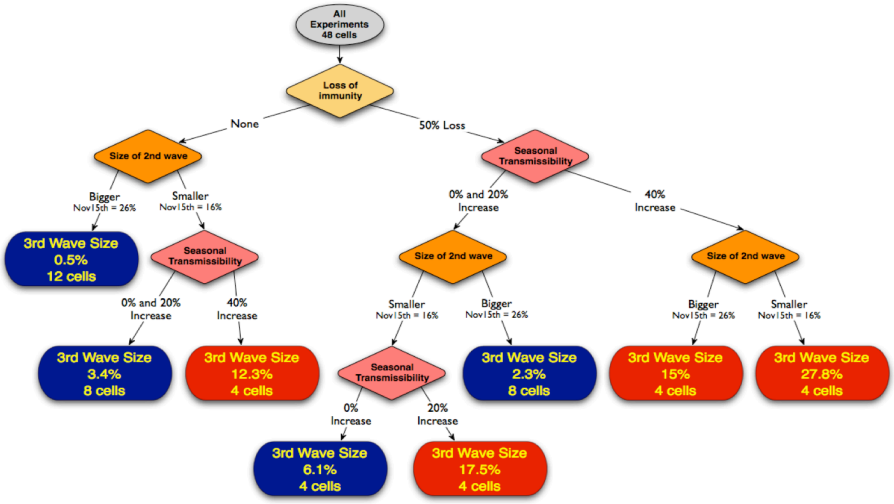

VIRGINIA BIOINFORMATICS INSTITUTE AT VIRGINIA TECH
Network Dynamics and Simulation Science Laboratory
Virginia Tech

Conclusions

- Loss of immunity and changes in transmissibility due to Seasonal factors are needed to cause a large 3rd wave of infections
- Better situational awareness is needed as results depend on knowing how many people have been infected to date
- Results are insensitive to when these changes occur (now or much later have similar results)
- Transmission simulations can not predict how likely these conditions will arise

Appendix E. *in silico* Surveillance Website

Further information concerning Chapter II, including animations of the outbreak detection process, can be found at the following website: http://ndssl.vbi.vt.edu/insilicoSurveillance/index.html

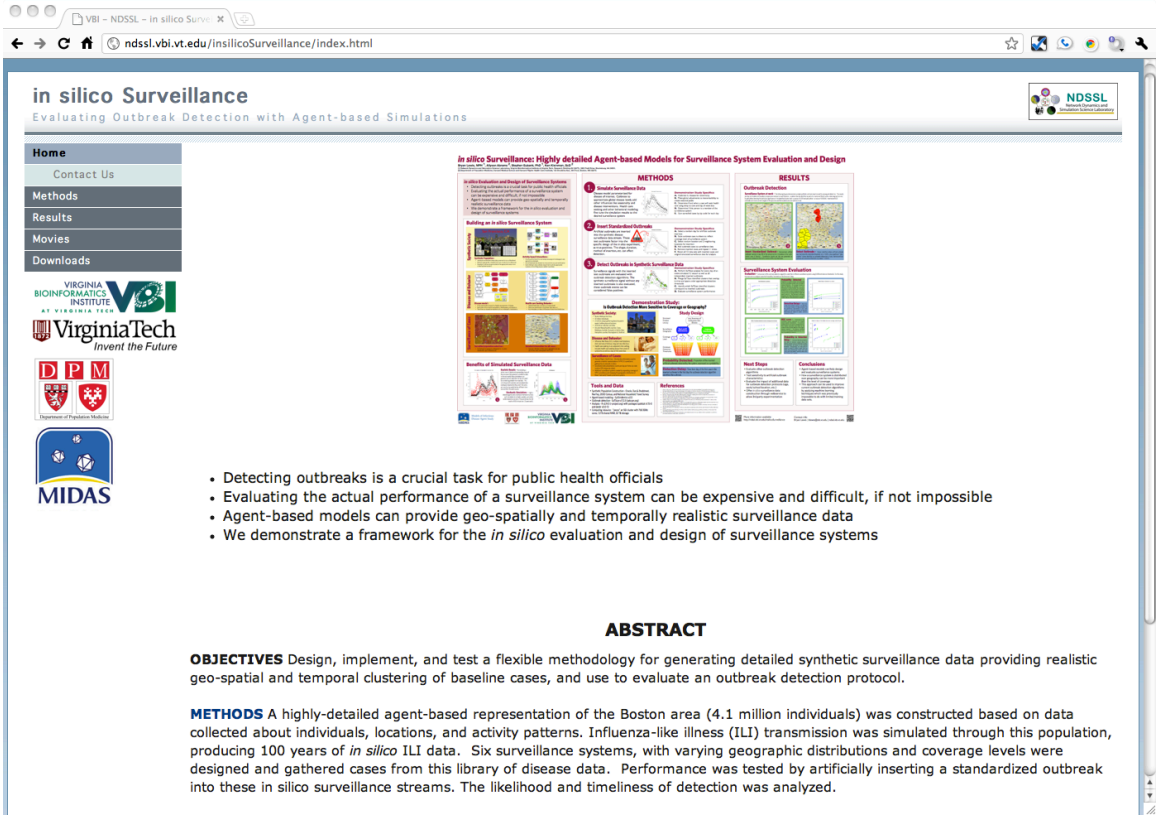

Appendix F. *in silico* Surveillance Poster

An alternative presentation of the material in Chapter II is a poster that was presented at the American Public Health Association annual convention in Denver in early November of 2010. A full-size PDF of this poster can be downloaded here: http://db.tt/5DzcSE7

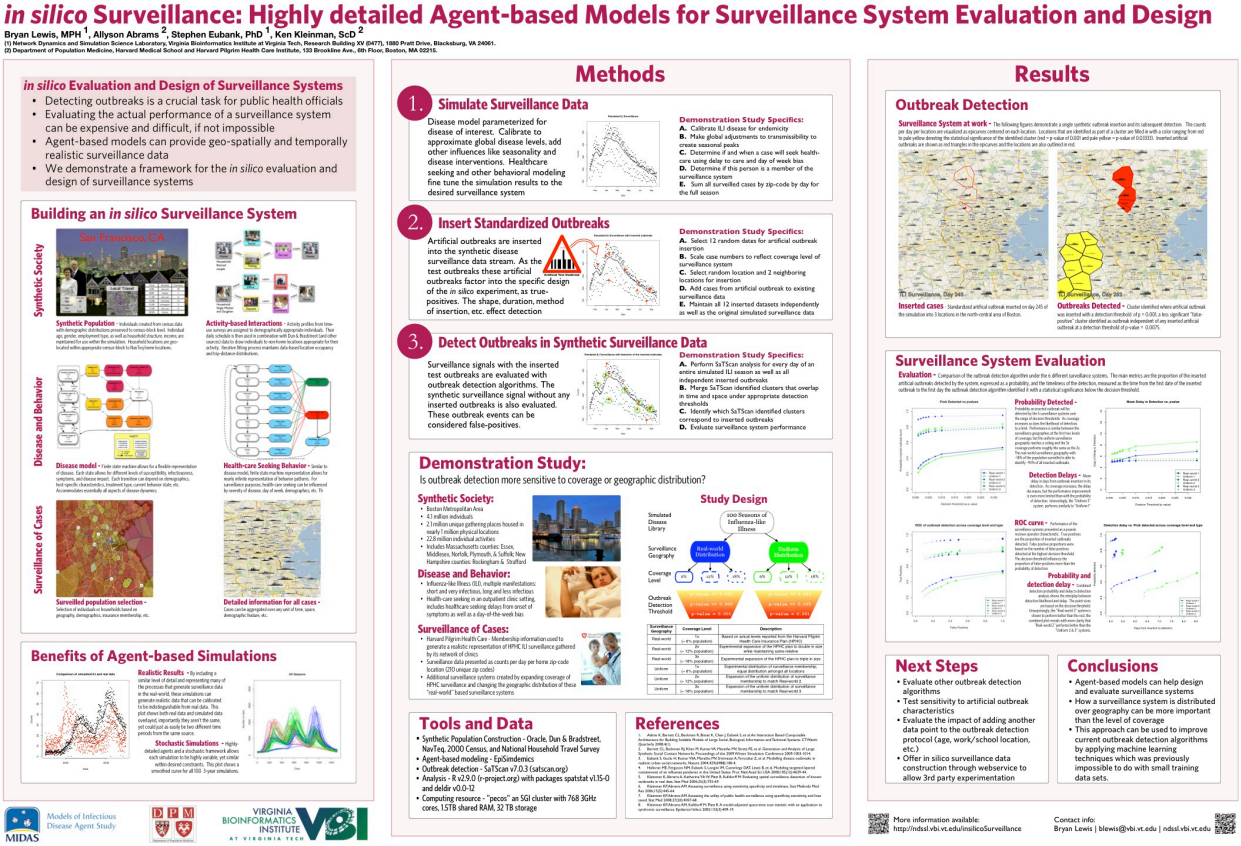

Appendix G. Epinome Overview

The software developed at the University of Utah is described and demonstrated in the on-line video hosted at:
http://www.sci.utah.edu/gallery2/v/software/epinome/

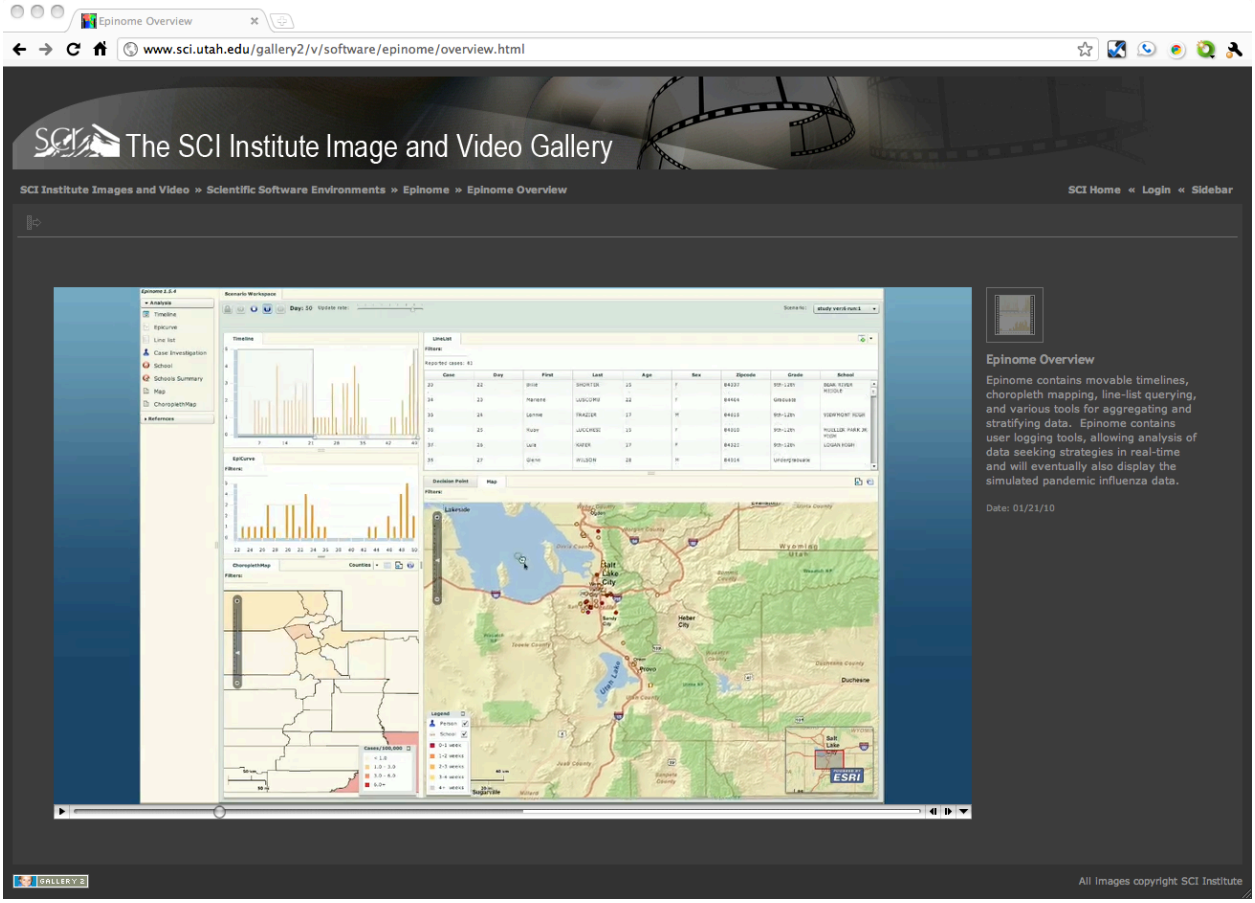

Appendix H. Annotated List of Figures and Tables

Introduction

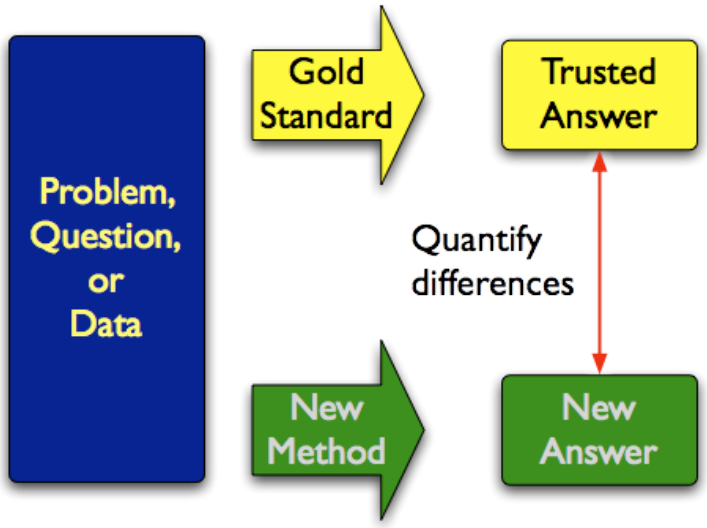

Figure 1. Diagram of process used to develop confidence in a new method

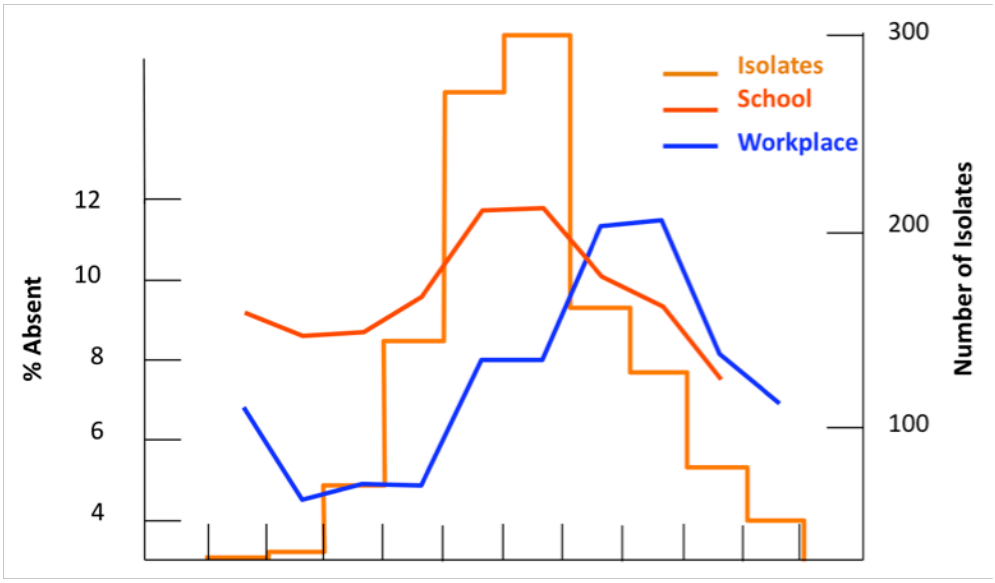

Figure 2. Real world observations from an influenza epidemic in Houston, Texas 1978. School-based absenteeism precedes workplace absenteeism as the epidemic (as measure in influenza positive isolates) peaks.

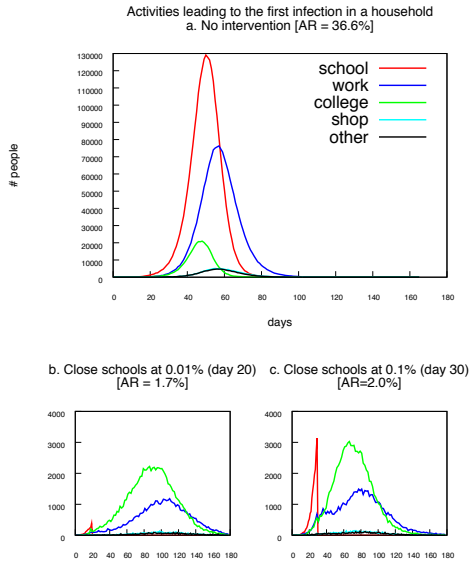

Figure 3. Infection counts of simulated influenza epidemics over time, separated by activity leading to the first infection in a household. A) The unmitigated epidemic infects 36.6% of the population B) Schools are closed 20 days into the epidemic when only 0.01% of the population is infected in a single day C) Schools are closed 30 days into the epidemic when 0.1% of the population is infected on a single day.

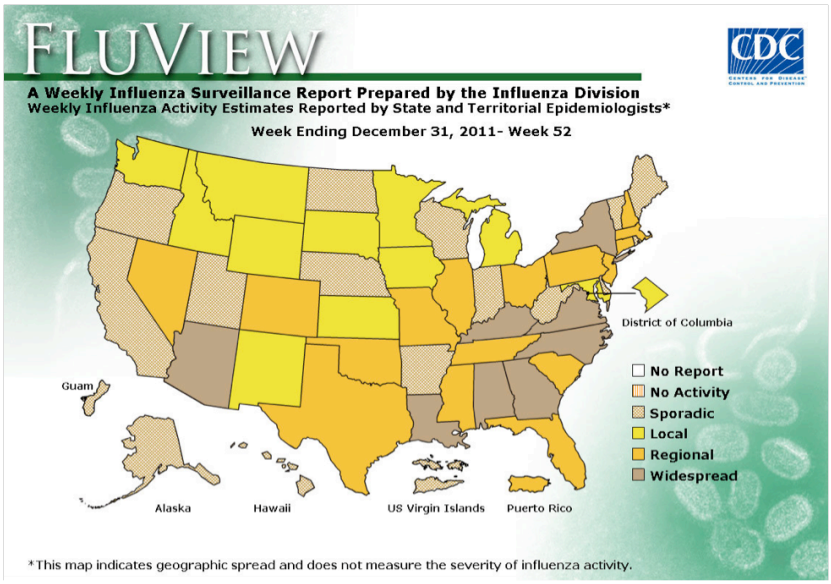

Figure 4. Weekly Influenza Surveillance Report, for December 24-31 (week 52), 2010 (available: http://www.cdc.gov/flu/weekly/usmap.htm; accessed Jan 2nd, 2011)

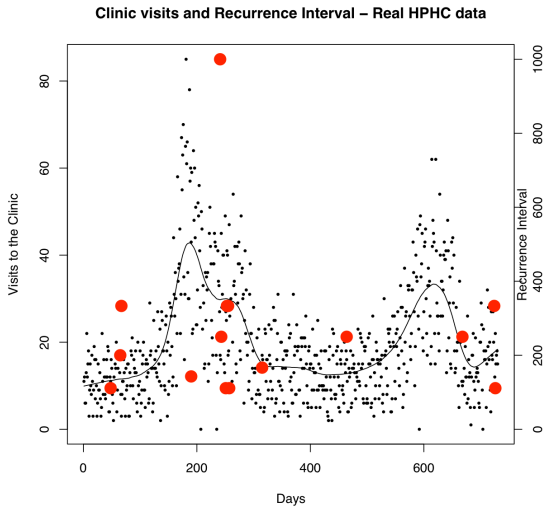

Figure 5. Detection of outbreaks for 2 years of Influenza-like Illness Surveillance data from the Harvard Pilgrim Health Care (HPHC) clinics in the Boston Metro Area. Black dots represent actual counts per day, black line is a smoothed fit of these counts, and the red dots indicate the identification of a cluster and its statistical significance (represented as a recurrence interval, 1/p-value)

Scenario% compliance/ ascertainment	Intervention threshold, %	R_0 = 1.9 (2.1)			R_0 = 2.4			R_0 = 3.0		
		Imperial	UW	VBI	Imperial	UW	VBI	Imperial	UW	VBI
1	NA	42.4	46.8	44.7	52.4	52.4	51.1	58.8	58.8	56.5
		(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)	(0)
2 30/60	1	7.3	2.8	3.9	15.5	4.1	9.7	27.4	8.5	20.1
		(104)	(38)	(59.1)	(239)	(61.2)	(140.4)	(421)	(138.4)	(275.1)
3 60/60	0.1	1.1	0.31	1.3	4.1	0.41	3.2	18.1	1.03	9.4
		(17.9)	(4.3)	(29.7)	(69.8)	(6.2)	(70.6)	(300)	(17.1)	(189)
4 60/80	0.01	0.22	0.04	0.10	0.66	0.05	1.4	10.1	0.12	3.8
		(4.6)	0.76)	(3.2)	(14.2)	(1.1)	(41.7)	(210)	(2.5)	(105.9)
5 90/60	0.1	0.17	0.30	1.2	0.66	0.30	2.7	11.5	0.86	6.4
		(3.2)	(4.1)	(28.8)	(14.2)	(5.8)	(59.4)	(192)	(14.2)	(132.2)
6 90/80	0.01	0.20	0.04	0.07	0.54	0.05	1.2	5.4	0.11	2.8
		(4.2)	(0.76)	(2.3)	(9.4)	(1.0)	(35.5)	(112)	(2.1)	(80.2)

The Imperial/Pitt model results are based on an average of 10 realizations, the UW/LANL results on an average of 5 realizations, and the VBI results mostly on one realization.

Table 1. The combined scenarios of targeted layered containment

	Scenario (Compliance %/Ascertainment %)										
	1 Base	2 30/60		3 60/60		4 60/80		5 90/60		6 90/80	
Intervention	case	Achieved	Compliance	Achieved	Compliance	Achieved	Compliance	Achieved	Compliance	Achieved	Compliance
Symptomatic cases ascertained		60		60		80		60		80	
In ascertained case household											
Threshold	–	1.0		0.1		0.01		0.1		0.01	
Index case treated	–	100	*	100		100		100		100	
Contacts prophylaxed (TAP)	–	100	*	100	*	100	*	100	*	100	*
Home isolation of cases	†		60		60		60		90		90
Quarantine of contacts	–		30		60		60		90		90
School closure	–	100		100		100		100		100	
Threshold	–	1.0		0.1		0.01		0.1		0.01	
Children kept home‡	–		30		60		60		90		90
Workplace distancing	–	50		50		50		50		50	
Threshold	–	1.0		0.1		0.01		0.1		0.01	
Liberal leave	†	100		100		100		100		100	
Community social distancing	–	50		50		50		50		50	
Threshold	–	1.0		0.1		0.01		0.1		0.01	

All numerical values are percentages.
*UW/LANL model assumes 5% stop taking drug after 1 day.
†In all three models, a proportion of symptomatic people retire to home even without intervention.
‡Compliance is % reduction in contacts or contact probabilities outside home.

Table 2. Illness attack rates (%) and (antiviral courses per 1,000) using scenarios described in Table 1 in the Chicago population

	Scenario 1. No intervention			Scenario 2			Scenario 3		
	Imperial	UW	VBI	Imperial	UW	VBI	Imperial	UW	VBI
Illness attack rates	42.4	46.8	44.7	7.3	2.8	3.9	1.1	0.31	1.3
Places									
Home	33.1	39.4	41.1	48.3	58	45.9	50.4	59	36.9
Work	21.8	14.5	28.6	12.9	10	27.8	13.5	10	18.7
School	16.0	18.8	23.3	11.7	11	9.6	9.0	11	2.7
Day care	–	1.1	–	–	0	–	–	0	–
Play group	–	0.8	–	–	0	–	–	0	–
College	–	–	3.3	–	–	12.3	–	–	40.0
Shopping	–	–	2.0	–	–	2.4	–	–	1.0
Neighborhood	–	17.7	–	–	15	–	–	15	–
Neighborhood clusters	–	7.7	–	–	5	–	–	4	–
Other/Community	29.0	0	1.7	26.6	0	2.0	23.8	0	0.8
Totals									
Primary Groups*	70.9	72.7	93.0	72.9	79	83.3	72.9	79	58.3
Community†	29.0	25.4	3.7	26.6	20	4.4	23.8	19	1.8

*Includes home, school, workplace, and for the UW/LANL model, day care and play groups.
†Includes groups subject to community social distancing.

Table 3. Percentage of infections by place and scenario, R0 1.9 (2.1) in the Chicago population

	Illness (infection)	Attack rate, %
Scenario	UW/LANL	Imperial/Pitt
No intervention	47 (70)	42 (63)
Social distancing*	39 (58)	–
Partial scenario 2†	23 (35)	–
Full TLC (scenario 2)‡	0.13 (0.20)	0.30 (0.45)

Threshold is an illness attack rate of 1/1,000 nationally for all interventions except school closure. School closure is implemented locally at the local threshold of 1/1,000 illness attack rate. Otherwise similar to scenario 2 (30/60) when applicable. UW/LANL model $R_0 = 2.1$; Imperial/Pitt model $R_0 = 1.9$.
*Only 50% community social distancing and 50% reduction in long distance travel, nothing else.
†Scenario 2, 50% reduction in long distance travel; but no TAP, treatment only, no school closure, no liberal leave.
‡Scenario 2, school closure at local threshold; 50% reduction in long-distance travel.

Table 4. U.S. national illness (infection) attack rates using three national intervention strategies in the U.S. population models

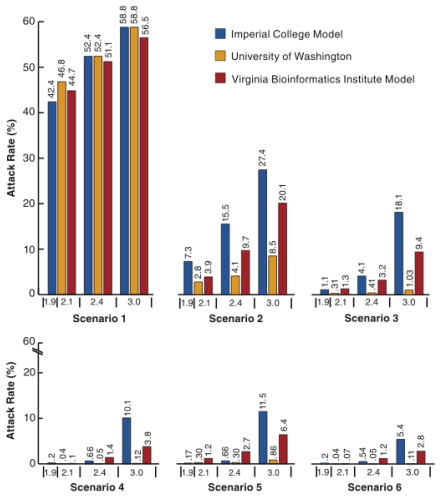

Fig. 1. Influenza illness attack rates for three R_0 values without intervention and with five scenarios of TLC intervention by using the three different models (Chicago population). See Table 1 for a description of scenarios. The R_0 values of 1.9 and 2.1 are considered as a single comparison.

Figure 1. Influenza illness attack rates for three R0 values without intervention and with five scenarios of TLC intervention by using the three different models (Chicago population). See Table 1 for a description of scenarios. The R0 values of 1.9 and 2.1 are considered as a single comparison.

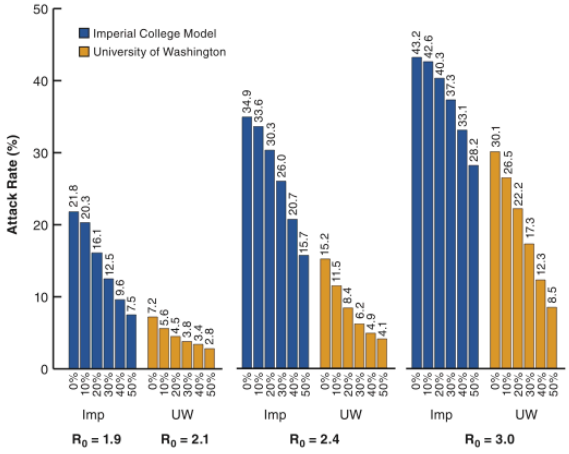

Fig. 2. Sensitivity analysis for workplace and community social distancing. Scenario 2, with community and workplace social distancing being varied between 0% and 50%, and three R_0 values (Chicago population). Only the UW/LANL and Imperial/Pitt models were used. The VBI model is insensitive to changes in this aspect of community social distancing.

Figure 2. Sensitivity analysis for workplace and community social distancing. Scenario 2, with community and workplace social distancing being varied between 0% and 50%, and three R0 values (Chicago population). Only the UW/LANL and Imperial/Pitt models were used. The VBI model is insensitive to changes in this aspect of community social distancing.

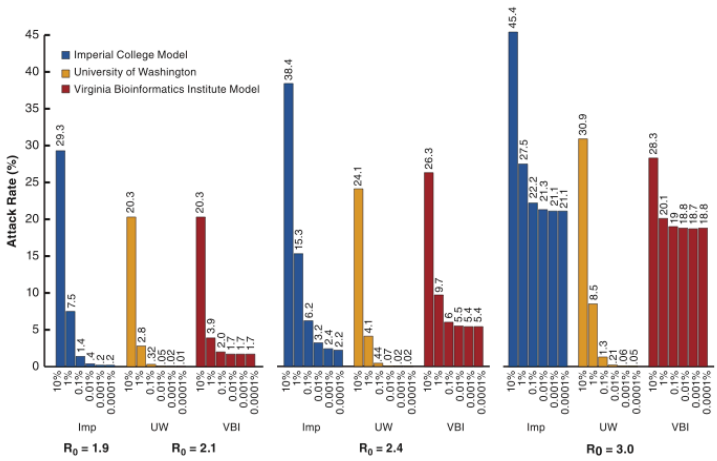

Fig. 3. Sensitivity to changing thresholds for all interventions simultaneously for the three models. Scenario 2 and three R_0 values, with threshold for triggering all measures being varied between 10% and 0.0001% cumulative illness attack rates. Chicago population.

Figure 3. Sensitivity to changing thresholds for all in- terventions simultaneously for the three models. Sce- nario 2 and three R0 values, with threshold for trigger- ing all measures being varied between 10% and 0.0001% cumulative illness attack rates. Chicago population.

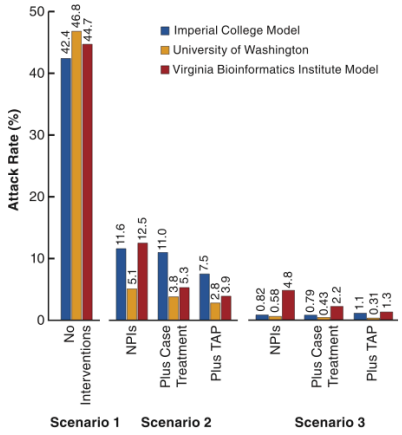

Fig. 4. Comparison of no intervention with intervention scenarios 2 and 3 using just NPIs, NPI with addition of just treatment of ascertained cases (Plus Case Treatment), and NPI with addition of treatment of ascertained cases and targeted antiviral prophylaxis (Plus TAP) of their household contacts. Scenario 1: no intervention; scenario 2: just NPI, with treatment only; with TAP (base case scenario 2); scenario 3: just NPI, with treatment only; with TAP and treatment (base case scenario 3); R_0 of 1.9 (2.1). Chicago population.

Figure 4. Comparison of no intervention with intervention scenarios 2 and 3 using just NPIs, NPI with addition of just treatment of ascertained cases (Plus Case Treatment), and NPI with addition of treatment of ascertained cases and targeted antiviral prophylaxis (Plus TAP) of their household contacts. Scenario 1: no intervention; scenario 2: just NPI, with treatment only; with TAP (base case scenario 2); scenario 3: just NPI, with treatment only; with TAP and treatment (base case scenario 3); R0 of 1.9 (2.1). Chicago population.

Chapter II

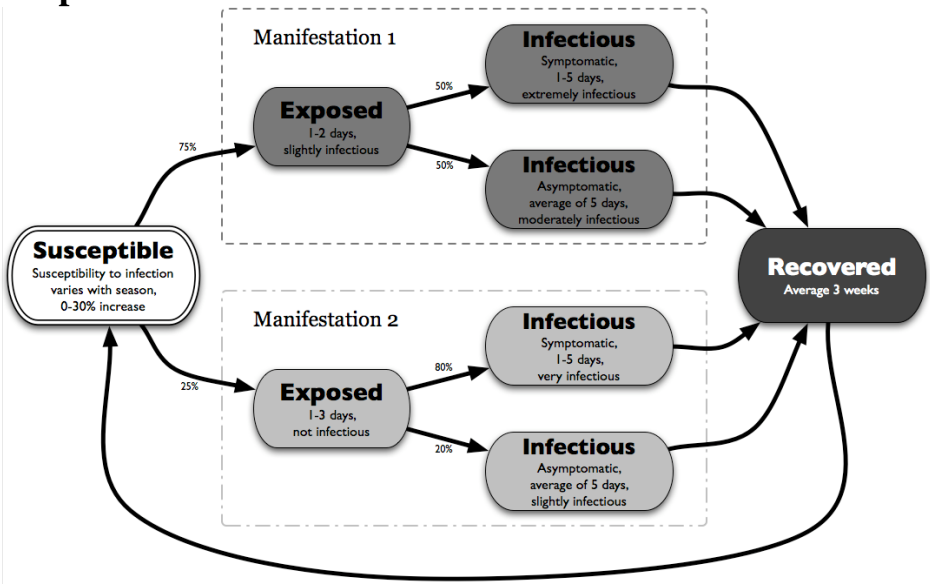

Figure 1. Finite state machine representation allows for a flexible representation of a disease process. Each state determines the duration in that state, level of symptoms, susceptibility, and infectiousness. For example, susceptible individuals have a 75% chance that upon exposure to an infectious contact and successful infection they will transition to the latent Exposed state in ILI Manifestation 1; after 1 to 2 days they will then transition to one of the more Infectious states with a 50% chance of having symptoms.

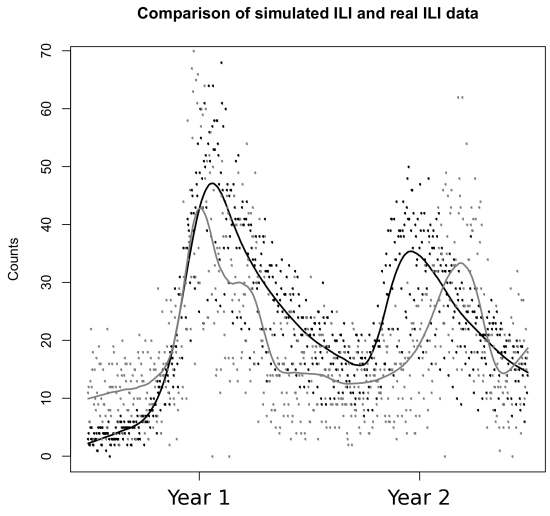

Figure 2. Two years of real ILI data compared to 2 years of data from a single simulation. Each dot represents the total number of cases seen in the surveillance stream on a single day; the curves are a smoothed fit to these data.

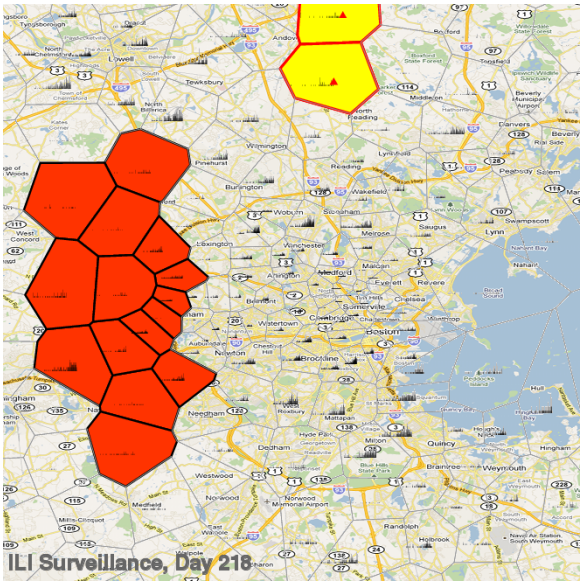

Figure 3. Simulated ILI surveillance data for downtown Boston as captured by the "Natural 6" surveillance system. Surveillance counts per day centered in each zip code location are shown as histograms within each zip code. Detection of an inserted test outbreak (red triangle) is indicated by red bordered zip codes and a false-positive outbreak by black bordered zip codes.

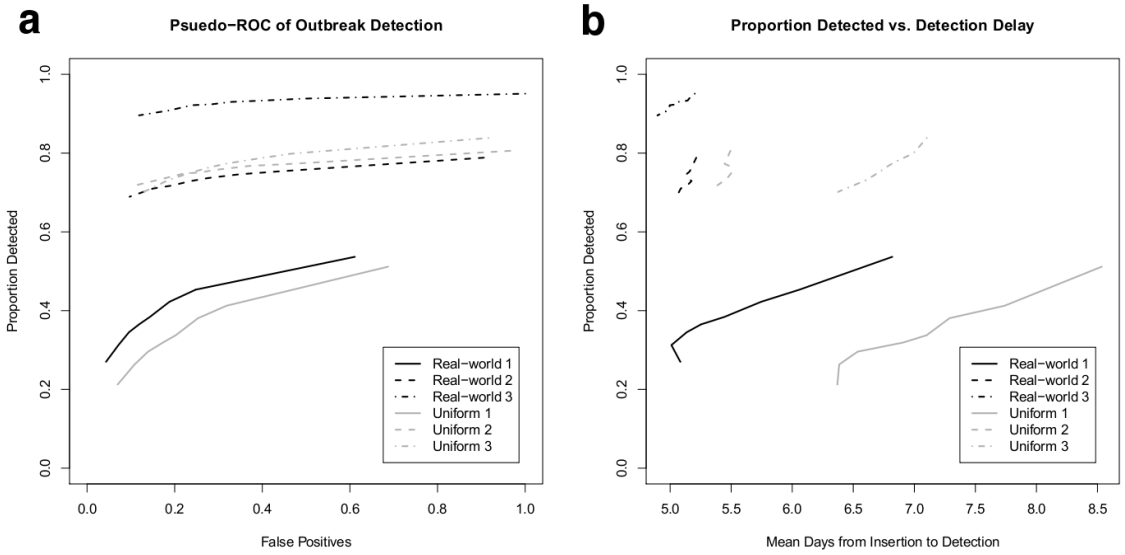

Figure 4. A) Pseudo-ROC curves of outbreak detection. Proportion detected for each surveillance system vs. proportion of all false-positives identified. B) Proportion detected vs. mean days to detection across decision thresholds.

#	Policy	Contacts included in intervention	Actions performed on intervened contacts
1	Least intense	None	None
2	Status quo – Maintain current Utah state policy	High-risk contacts only: - Less than 1 year - Under-immunized (incomplete course of vaccines) - High-risk co-morbidity (pregnancy, immuno-compromised, etc.)	- Report all symptomatic cases - Prophylax all uninfected - Quarantine children under 8 years old - Immunize any under-immunized and uninfected individuals
3	Most intense	All contacts of more than an hour duration	Same as policy #2

Table 1. Statewide policy choices for pertussis control, in terms of a package of interventions and their targets.

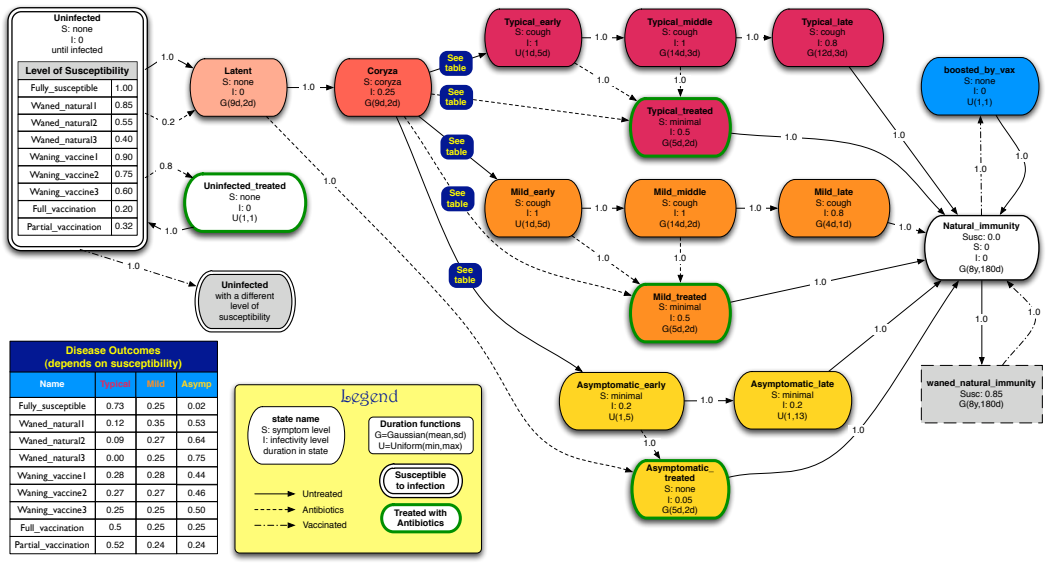

Figure 1. A finite state machine representation of the pertussis disease model, all states, dwell times, transition probabilities, relative level of susceptibility and infectivity, are explicitly represented.

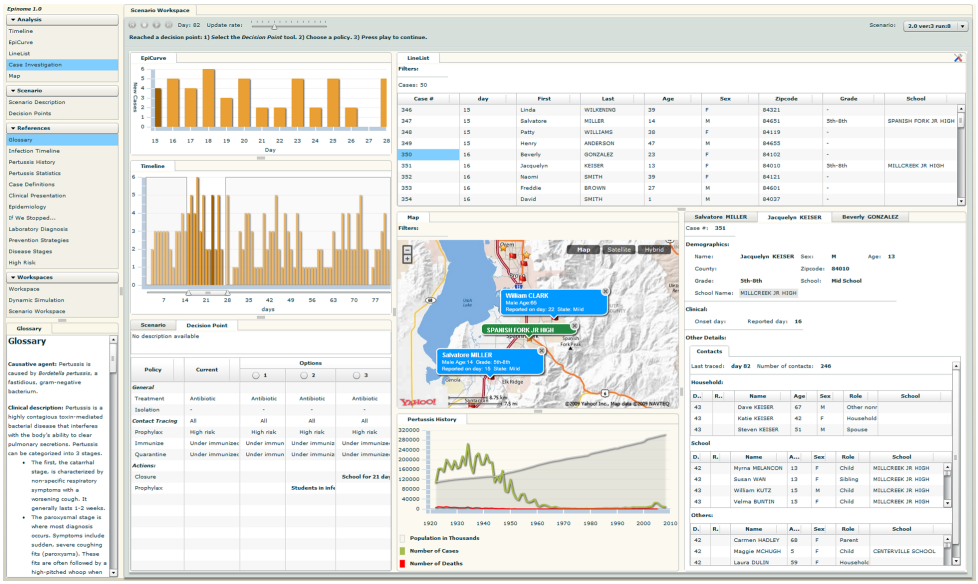

Figure 2. Screenshot of Epinome representing pertussis disease transmission in Utah.

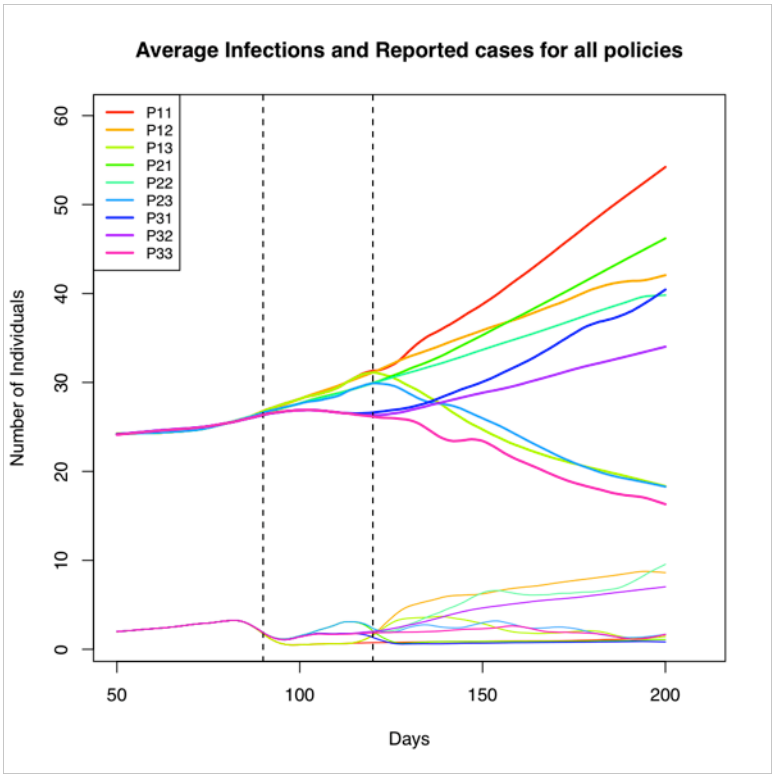

Figure 3. Total daily infections (thick lines) and reported cases (thin lines) for all policy choices averaged over all 20 iterations.

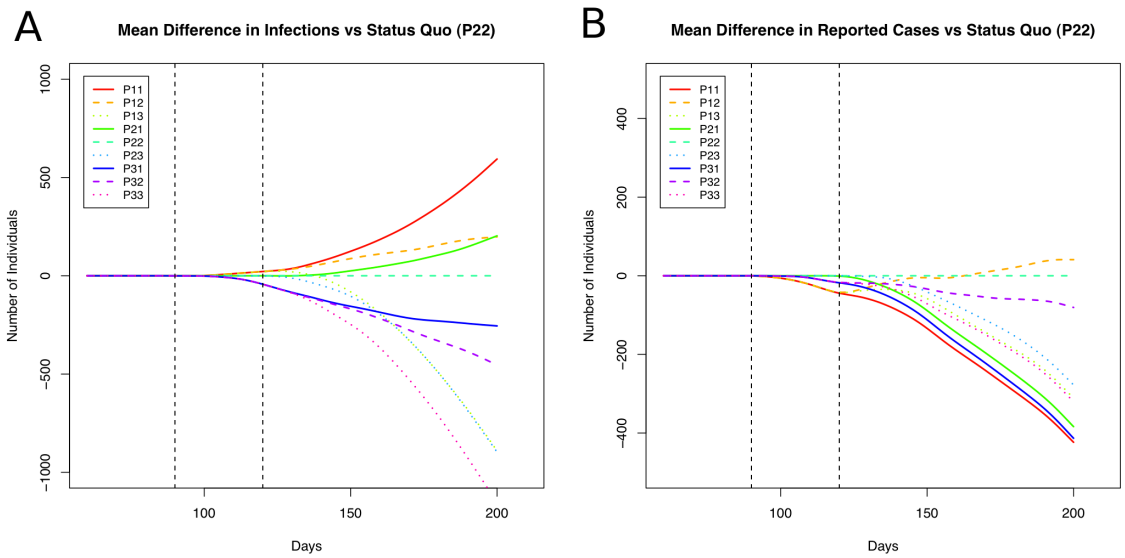

Figure 4. Outcome comparison across policies, cumulative mean differences from the moderate intensity policy (P22). A) The difference in infections (ground truth) B) reported cases

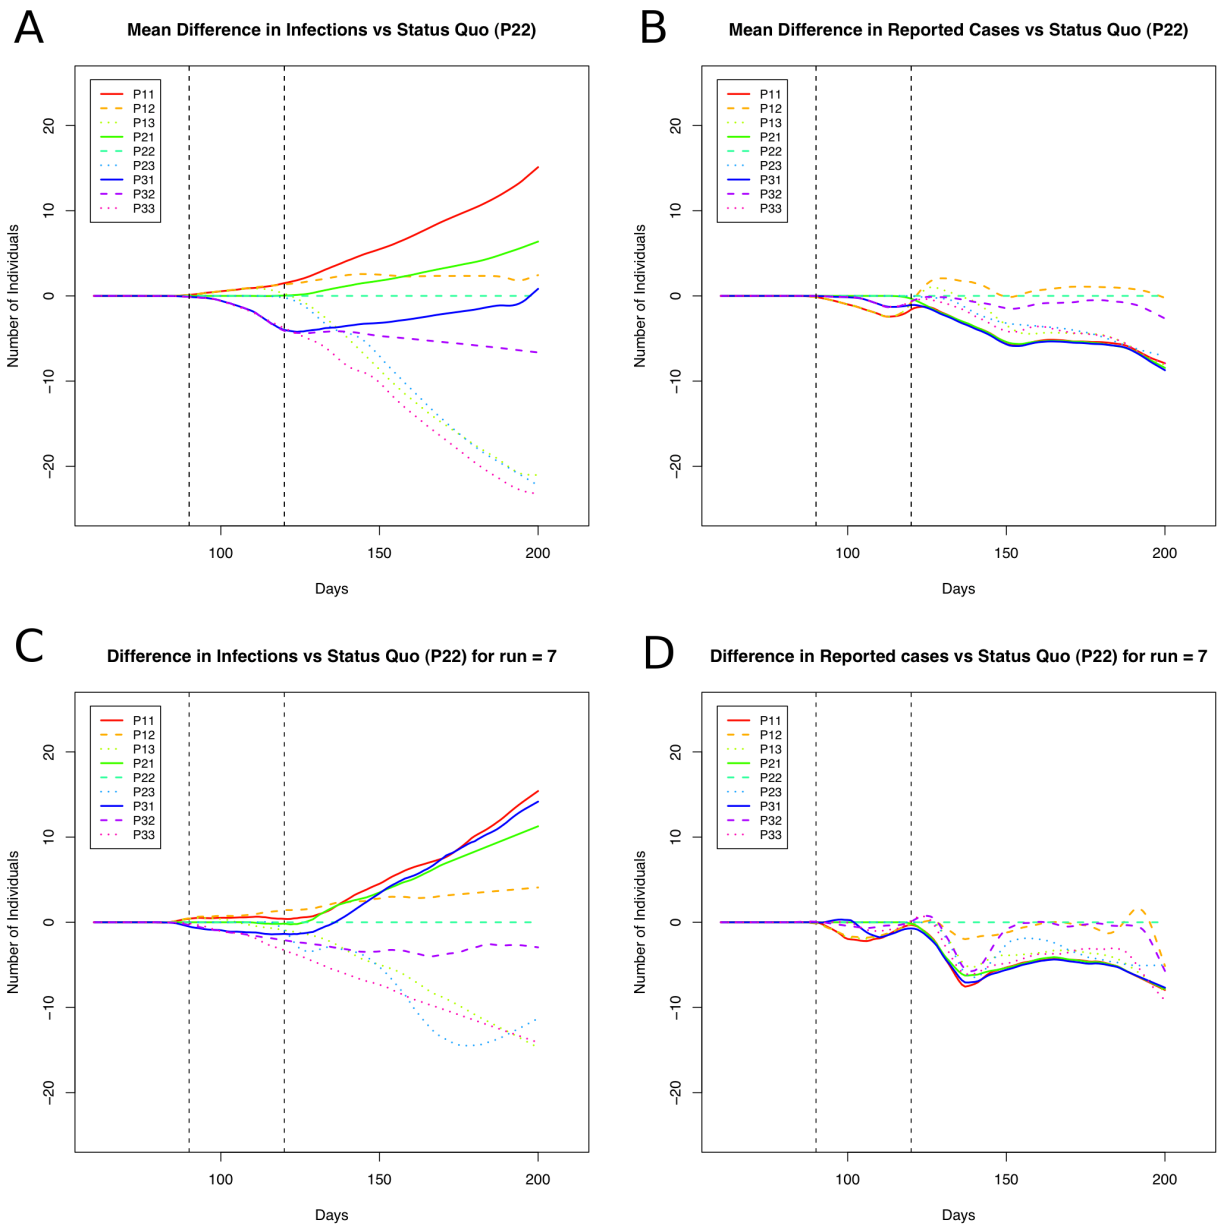

Figure 5. Comparison between average over 20 iterations and an individual stochastic realization (run 7) as a difference from a baseline of the status quo policy (P22). A) Mean infection counts per day B) Mean reported case counts per day C) Single run, infection counts per day D) Single run, reported case counts per day

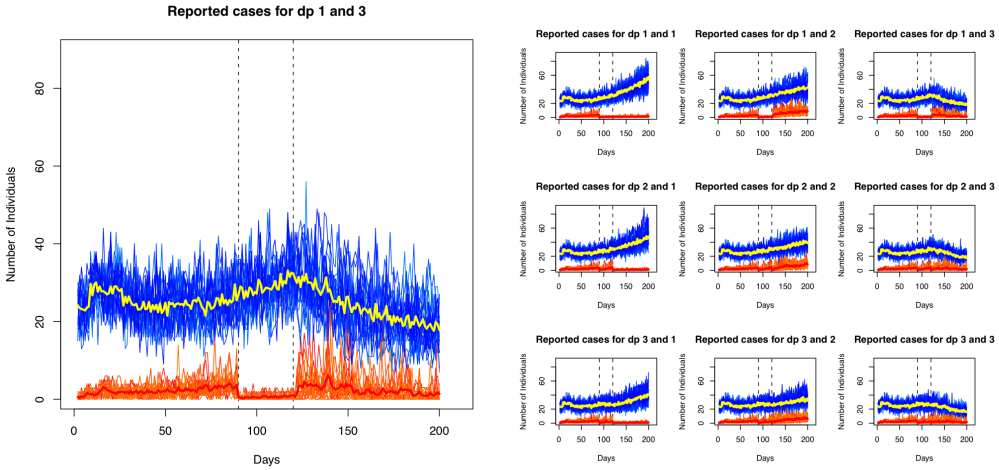

Figure 6. Number of pertussis infections (blue) and number of reported cases (red-orange) per day of simulation. A.) All 20 stochastic realizations for a single set of decisions (trace no contacts at day 90 and trace all contacts at day 120) B.) Same plot but for all 9 combinations of policy decisions.