The Inter-Region Epidemic Dynamics Model

February 7th, 2011

Jon Parker
National Center for the Study of Preparedness and Catastrophic Event Response (PACER)
The John Hopkins University
Senior Software Engineer
5801 Smith Ave
Davis Building, Suite 3200
Baltimore MD 21209
410-735-6400
jparker5@jhmi.edu

Abstract

The Inter-Region Epidemic Dynamics (IRED) Model efficiently simulates the spread of contagious disease throughout a large multi-patch region like the United States. For any initial geographic distribution of index cases, the model quickly generates an ensemble of stochastic realizations and computes the mean and standard deviation of incidence for every patch. Spatially heterogeneous epidemic severity can thus be assessed, with relevant estimates available for local contingency planning. Specifically, the model can estimate whether an epidemic beginning in one patch will spread to other patches in a statistically reliable fashion; thus helping to determine when mathematically rigorous epidemic forecasts can be made. The model can be used to study the impact of government-imposed travel restrictions or endogenous changes in travel behavior. Finally, it is capable of addressing comparative questions: Which generates a larger epidemic: 50 initial cases all in Chicago or 15 initial cases in Chicago plus 15 in Boston?

Keywords: Epidemic Simulation, Forecast, Travel Patterns.

The Inter-Region Epidemic Dynamics Model

Introduction

Modern mathematical epidemiology began with the Kermack-McKendrick model of 1927 [14]. This simple elegant system of differential equations posited perfect mixing within a single population. The model allows individuals to transit from the susceptible pool, or "compartment," to the infected state and finally to the removed (recovered or dead) one. Despite its simplicity, the Kermack-McKendrick model brought the threshold nature of epidemics sharply into focus and clearly explained "herd immunity," wherein immunity of a subpopulation can make outbreaks fizzle, protecting the remaining herd. These seminal equations spawned a voluminous and valuable scientific literature extending the basic model in many directions [16, 17]. Here we are interested in two of these.

First, the initial model was deterministic. NT Bailey [1] published a stochastic generalization in 1953. Whittle [2], and then many others, deepened the stochastic line of work [6, 7, 8]. The original equations also did not include space. Another large literature introduces multiple patches and the epidemiological effect of movement among them [3, 7, 8, 13]. There is also a smaller literature combining these, to produce stochastic multi-patch (also called meta-population) models [12]. The model introduced here contributes to this rather rarified literature an efficient and practical tool. The Inter-Region Epidemic Dynamics (IRED) Model is a stochastic multi-patch model that easily and efficiently computes the mean disease incidence and the standard deviation of disease incidence, for every patch in the model, given any initial distribution of index cases. Hence, it allows one to generate confidence intervals for the average level of disease in every patch.

A challenge faced when implementing a stochastic multi-patch model is to determine just how tightly or loosely any two patches should be coupled. Longini & Rvachev and Epstein et.al used international flight data to connect the patches in their models [21, 22]. This is realistic, but cumbersome and data-intensive. However, Ferguson et.al and Parker & Epstein use a gravity model to specify the patch-to-patch travel behavior that takes place in their epidemic models [18, 19, 20]. This is econometrically defensible and far more flexible and efficient.

From a policy standpoint, it is important to know *how robust are epidemic model outputs to changes in the patch-to-patch travel patterns encoded in these gravity models?* Little such detailed sensitivity analysis appears in the literature. The Inter-Region Epidemic Dynamics (IRED) Model is designed to fill this important gap. The IRED Model is well suited to this task due to its efficient computational design. Not only do individual stochastic IRED simulations proceed quickly, but the results from many simulations are automatically gathered and analyzed.

The IRED Model is also a highly capable modeling platform able to address other interesting epidemiologic questions. Some questions examined later in this paper are: How predictable are inter-region epidemic dynamics if only, say, 25 people are sick at the beginning of an epidemic? How does predictability change if 50 people are sick to begin with? Is short-term epidemic forecasting possible in either case? If so, how reliable are those forecasts X days into the future? How effective is implementing travel restrictions Y days into an epidemic? Finally, the IRED Model can facilitate comparisons involving complex spatial initial conditions: for example, we can quickly compare an epidemic with 50 initial cases all in Los Angeles to an epidemic with 15 initial cases in Los Angeles and an additional 15 cases in Boston.

Discussion

The general problem has two aspects. We need to determine how likely infections are to spread from one region to another in a multi-region epidemic model (the regions/patches need not be contiguous). We also need to determine how *robust* our forecasts are to our coupling assumptions. When both of these issues are addressed a wide variety of more interesting epidemiologic questions can be tackled.

Methodology: The IRED Model

Clearly, inter-region spread rates are determined by people's travel behavior. Unfortunately, generating a single travel matrix that reflects real-world data on all modes of travel is extremely difficult to do. Creating such a matrix requires collecting, cleaning, and properly combining data on driving behavior, commuter rail schedules, airline flight traffic, and possibly other modes of travel.

A way to relieve this data challenge was proposed by Voorhees, a traffic engineer, in 1956 [15]. He suggested using a formula that could plausibly predict the flow (i.e., number of trips) between any two regions. Voorhees referred to his formula as a gravity model due to its algebraic resemblance to Newton's Law of Gravitation. More formally, let F be a matrix that contains the estimated flow between different regions in a model. For example, if the entry F_{ij} = 20 then it is assumed that 20 people start trips in region i and end their trips upon arrival in region j. The entries in F are computed as follows:

$$F_{ij} = \frac{\left(p_i \cdot p_j\right)^a}{cost_{ij}^b} \quad (1)$$

Here $cost_{ij}$ is an estimate of the travel cost for trips between regions *i* and *j* (distance is usually a good proxy for cost), the *p*'s are the populations of the respective regions and *a* and *b* are tuning parameters. When a = 1, and b = 2, this is an inverse-square relation analogous to gravity. For an epidemiological application, see [4].

Using a gravity model to estimate travel patterns allows inter-region models to include a significantly higher number of regions. Without this widely applicable formula, computing the *F* matrix becomes an formidable task since the number of entries in *F* grows quadratically with the number of regions in the model. The IRED Model treats the *F* matrix as an input and uses it to generate other useful matrices. These matrices are then used to execute a series of stochastic events.

Before discussing the specifics of these stochastic events let us introduce the specifics of the Model. The IRED model posits a set of *N* places where each place has population p_i. A travel matrix *P* is given (or computed), wherein matrix entry P_{ij} is the probability that a person from region *i* travels to region *j* given that a trip occurs. An iteration of the IRED Model consists of these 4 steps:

0. Every person starts in their home region
1. Every person instantaneously travels to a random destination
2. Every person interacts with *m* people while in their destination
3. Every person instantaneously goes home
4. Every person transitions disease state if appropriate

These steps enable an SEIR (Susceptible, Exposed, Infected, and Recovered) epidemic to spread throughout the model population. People are moved from the susceptible state to the exposed state via "successful" contact with an infected person. At the end of each model iteration infected persons are moved to the removed state and exposed persons are moved to the infected state (further discussion below). To properly specify an IRED Model simulation one needs to input: a list of regions (and their populations), the travel matrix *P* (or the flows matrix F), the number of contacts per day *m*, and the number of susceptible, exposed, infectious, and recovered people in each region.

The IRED Model is computationally efficient because it also computes an interaction matrix *IM* that enables it to quickly account for the combined effect of everyone's travel behavior. The first step toward computing *IM* is to compute *P* from *F*. The travel matrix is computed by dividing each entry in *F* by its corresponding row sum as shown here:

$$P_{ij} = \frac{F_{ij}}{\sum_i F_{ij}} \tag{2}$$

Using *P*, the IRED Model is now able to compute the interaction matrix *IM*. The entry IM_{ij} is the probability that a person currently in region *i* contacts a person from region *j*. Pseudo-code that generates *IM* is shown below (this pseudo-code makes use of MATLAB function names).

POPS = Region populations //a row vector

TRAVELERS = diag(POPS) * P //a matrix

TOTALS = POPS * P //a row vector of column sums

$INTERACT_{ij}$= $TRAVELERS_{ij}$ / $TOTALS_j$ //a matrix

Another efficiency in the IRED Model is the fact that only infectious people need to have their itineraries executed. There is no reason to simulate the actions of uninfected people because the contacts they take part in cannot generate new infections – unless of course they include a "tracked" infectious person. To simulate the daily regime of an infected person from region *i*, the IRED Model first draws a random destination place using the i^{th} row of the *P* matrix. For the sake of explanation, assume region *k* was the selected destination. Now, the IRED Model uses the k^{th} row of *IM* to determine where a "contacted" person is from. The IRED Model then selects a random person from this region and executes a contact between these two agents. The k^{th} row of *IM* is used in this manner *m* times to yield the correct number of contacts per day for this specific infected person. Once all infected individuals have been operated on in this manner people can begin transitioning disease states (e.g. Exposed to Infected), as noted earlier.

The IRED framework is useful because all of its component operations can be optimized to execute extremely quickly. This results in a finished product that is capable of completing a national stochastic simulation in several seconds. The small execution time opens the door to a wide array of uses because it is possible to swiftly generate simulation data output.

IRED Model Application: Sensitivity Analysis

Recall, the gravity model shown in equation (1) has two parameters *a* and *b*. It is important to know how robust any conclusions will be to changes in these parameters. Here, we demonstrate that some measurements are rather robust while other measurements are rather

fragile. This dichotomy means some users could be willing to accept expedited modeling results while other users should be more cautions when interpreting expedited modeling results.

We begin by choosing 2 values for each *a* and *b*. This will enable us to execute 4 batches of 100 epidemic simulations where each batch has a different combination of parameters. All 400 simulations begins from the same initial conditions – 25 seed cases located in Los Angeles – and proceed to simulate a stochastic SEIR epidemic as it spread across the United States. The mean and variance of key statistics (measured on day 10 and 20) will then be extracted from each batch of results. The statistics gathered are: number of cases in the USA, number of cases in Los Angeles number of cases in Sacramento, and number of cases in Boston.

Table 1: Number of Cases on Day 10 and Day 20

		a = 1.0 , b = 2.0		a = 0.5 , b = 2.0		a = 1.0 , b = 1.5		a = 0.5 , b = 1.5	
		Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Day 10	U.S.A.	130	31	133	32	131	32	134	26
	Los Angeles	56	15	67	17	28	9.1	25	7.9
	Sacramento	1.2	1.6	.6	1	0.9	1.0	1.1	1.1
	Boston	0	0	0	0	0.4	0.8	0.1	0.4
Day 20	U.S.A.	986	254	994	247	1016	238	1024	206
	Los Angeles	224	60	323	85	77	20	59	16
	Sacramento	8	4.1	5.8	3.8	3.9	2.1	3.9	2.5
	Boston	.8	1.3	.3	.9	3.3	1.9	3.6	2.1

Table 1 simultaneously illustrates a few important points. First, the macroscopic statistic "Number of US Cases" is significantly more robust to different values of *a* and *b* than are the other more local statistics. Second, cities that are further away from the epicenter of the epidemic have statistics that are more sensitive to changes in parameters. Finally, the significance of parameter differences is smaller in the Day 10 results than it is in the Day 20 results.

This last point, that parameter differences are less important for measurements made a shorter time into the future, is important. It suggests that acceptable short term predictions can be made in the presence of some uncertainty about the parameters of the gravity model. This insight is explored further in the next section.

The sensitivity of various statistics to the sheer number of cases is also worth mentioning. The data in Table 2 shows how having a larger infectious population significantly reduces the index of dispersion for statistics measured in places that are likely to be within the epidemic's border.

Table 2: Number of Cases on Day 20 (a = 1, b = 2)

	25 initial cases in LA			50 initial cases in LA		
	Mean	Std Dev	σ^2/μ	Mean	Std Dev	σ^2/μ
US	987	253	64.9	1945	327	55.0
LA	224	60	16.1	436	68	10.6
Sacramento	8.0	4.1	2.1	14	4.5	1.4
Las Vegas	5.8	3.5	2.1	12	4.9	2.0

IRED Model Application: Short-term Forecasting

Tables 1 and 2 were compiled by executing multiple batches of 100 simulations. Every simulation began at time zero with all initially infectious people confined to Los Angeles. In the real world, nothing forces the seed cases to be located in a single region. Similarly, in a real world epidemic there is no "time equals zero" point. Accordingly, the IRED Model is capable of incorporating any set of initial conditions with a multinomial distribution of susceptible, exposed, infected, and recovered people. These arbitrary initial conditions can be used as the starting point from which to make near-term forecasts. In other words, given accurate initial conditions the IRED Model can be used as a “nowcasting” engine.

To demonstrate how the IRED Model can be used to make epidemic forecasts we first execute a single IRED Model simulation. The path of this single simulated epidemic is treated as the path of a *real* epidemic. The *real* epidemic begins with 50 initial cases all within Los Angeles. Next, 100 simulations with the same initial conditions, but different random seeds, are executed. We want to use the results from the batch of 100 simulations to predict the path of the unknown *real* epidemic. Figure 1 clearly shows how predicting Total US Cases at any given time will be difficult. Making accurate prediction before the epidemic peak is difficult because the index of dispersion (Variance in Num US Cases at time *t* divided by Mean US Cases at time *t*) is so high. Predictability increase dramatically (dispersion drops) when the epidemic is near its peak. However, at this point in an epidemic the value of an accurate prediction is questionable.

Figure 1: Index of Dispersion for Predicted US Cases (Red) and Total US Cases (Blue)

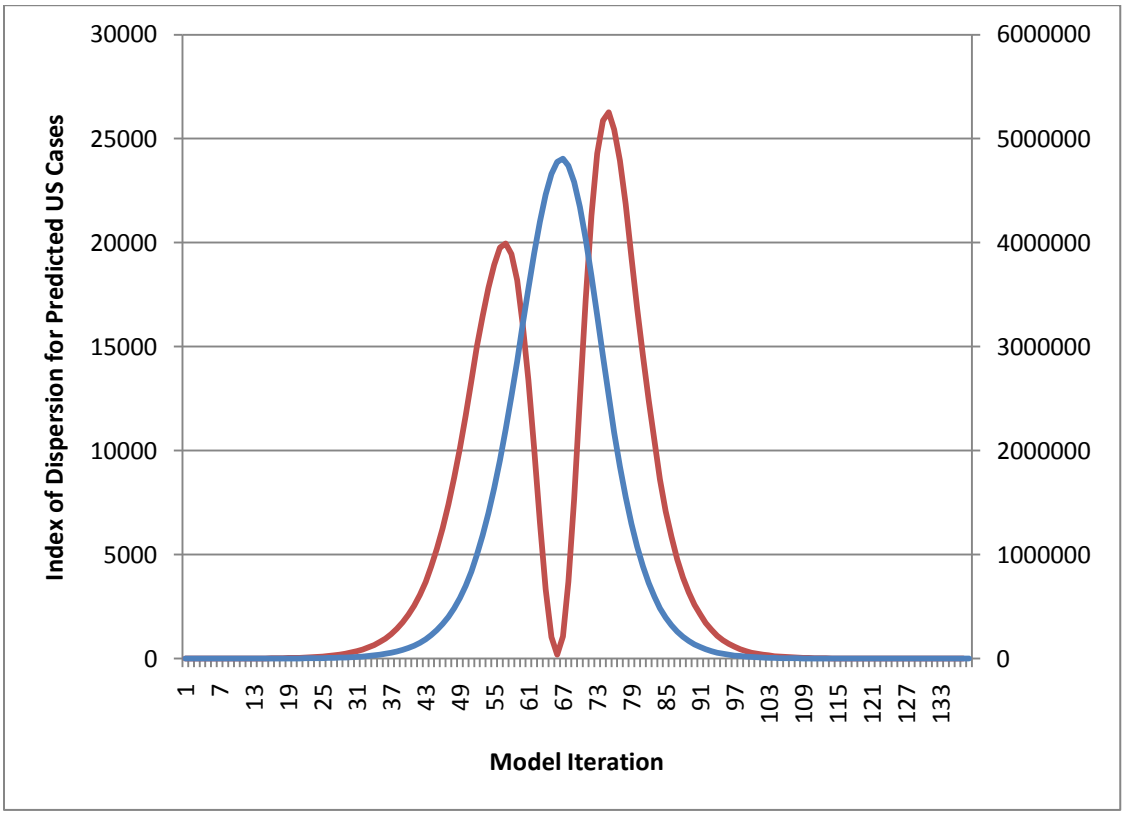

Making accurate epidemic predictions would be significantly easier if more than just the very initial conditions (50 cases in L.A.) could be used as input. With this in mind, we executed another batch of 100 simulated epidemics. Only this time the batch simulations were periodically synched to the *real* epidemic's path. This is analogous to loading the current weather, as collected by Doppler radar, into a weather model to enable better prediction. When this synchronizing approach is used the index of dispersion remains low enough to permit useful epidemic forecasting.

Figure 2: Index of Dispersion for Predicted US Cases (Red) and Total US Cases (Blue)

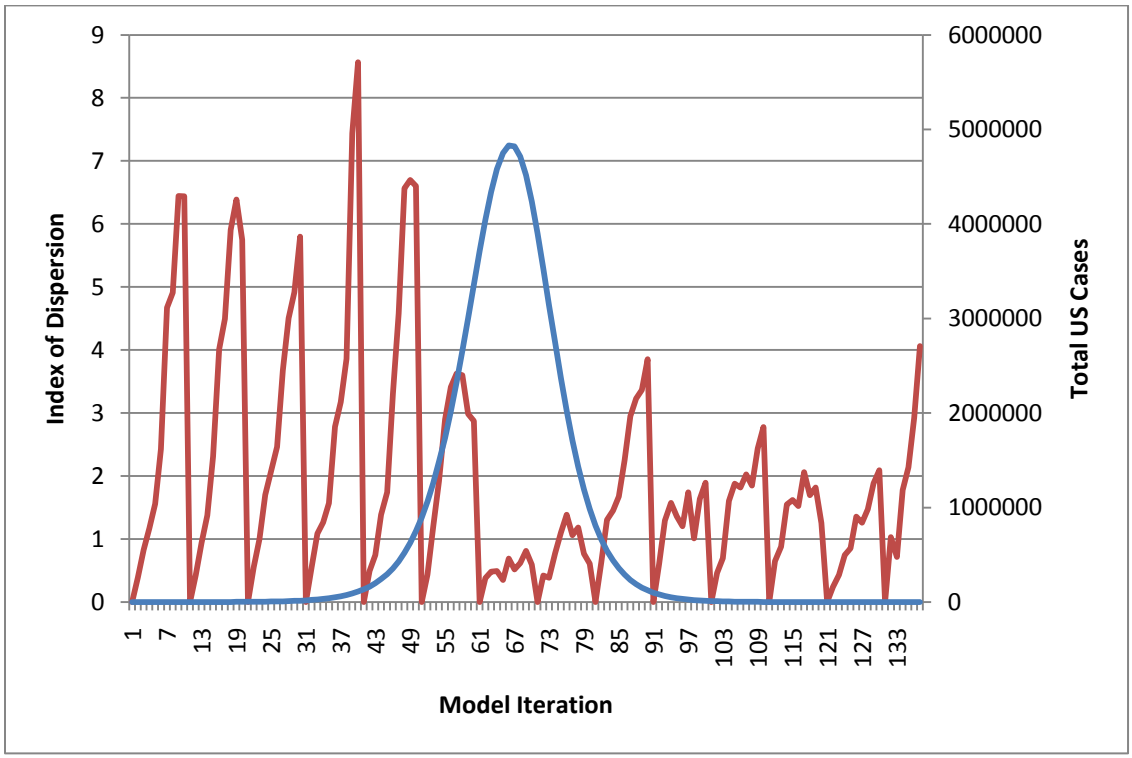

Challenges to Real World Short-term Forecasting

The prediction technique discussed above can, in principle, be used during real world epidemics. Enabling these types of predictions during real world epidemics requires making changes to the IRED Model's simple SEIR framework. It will also require the ability to gather

and input accurate, time-sensitive data on the number of people currently – and formerly – infected. Luckily, both hurdles may be less onerous then they appear.

Properly altering the IRED Model's simple SEIR structure to reflect a real world disease requires detailed information about that disease. The standard IRED Model assumes people are contagious for exactly one day. However, many pathogens cause people to be infectious for multiple days. This difference can be handled by inserting multiple infected states into the model. For example, the list of possible disease states can be extended from SEIR to $SEI_2I_1I_0R$. The index on the infected states denotes how many more days the person will be contagious. Other changes to the standard IRED Model can be made to support different types of pathogens. Highly contagious diseases can be modeled by increasing *m*, the number of people a person contacts each day or increasing the probability of disease transmission given a contact. Diseases with long incubation period can be modeled by introducing more Exposed states. Combining these adjustments might produce a state space like $SE_2E_1E_0I_2I_1I_0R$. It is likely that a set of simple changes like these could be used to closely mimic the real-world dynamics of many contagious diseases.

The data collection side of enabling short term epidemic forecasting is much less straight forward and seems daunting at first blush. However, this particular problem is likely to be diminished by the fact that the epidemic intensity in one region is typically related to the epidemic intensity of other spatially or temporally proximate regions.

Determining which values should be in a small collectable dataset to best enable accurate interpolation is a topic ripe for future study. Clearly, any interpolation scheme will benefit from additional data points. However, every data point will represent a real-world data collection effort; so adding more data points may be prohibitively costly. It would be interesting to study a

data collection scheme that supplied accurate data in different regions at different times. For example, collect NYC and Baltimore data during odd weeks and Boston and Philadelphia data during even weeks. A scheme like this may enable useful predictions because it will provide some current data for regions that are far away and some recently current data for regions close by. We suspect a data collection scheme like this will enable good prediction while keeping the number of simultaneously operating data collection sites to a minimum.

IRED Model Application: Travel Restriction Analysis

In this section we show how the impact of travel restrictions can be studied using the IRED Model. The IRED Model can simulate the impact of travel restriction by changing the *P* and *IM* matrices at a specific time in a given simulation. The restricted travel matrix prevents any region-to-region trips longer than 50 miles. Notice, this means the "wave of infection" can travel at most 100 miles in a day. The maximum travel distance is twice the travel restriction level because the *infector* and the *infectee* can both travel half the distance and meet in the middle. All simulations begin with 50 infected people in Los Angeles.

Figure 3: Impact of Travel Restrictions within the United States

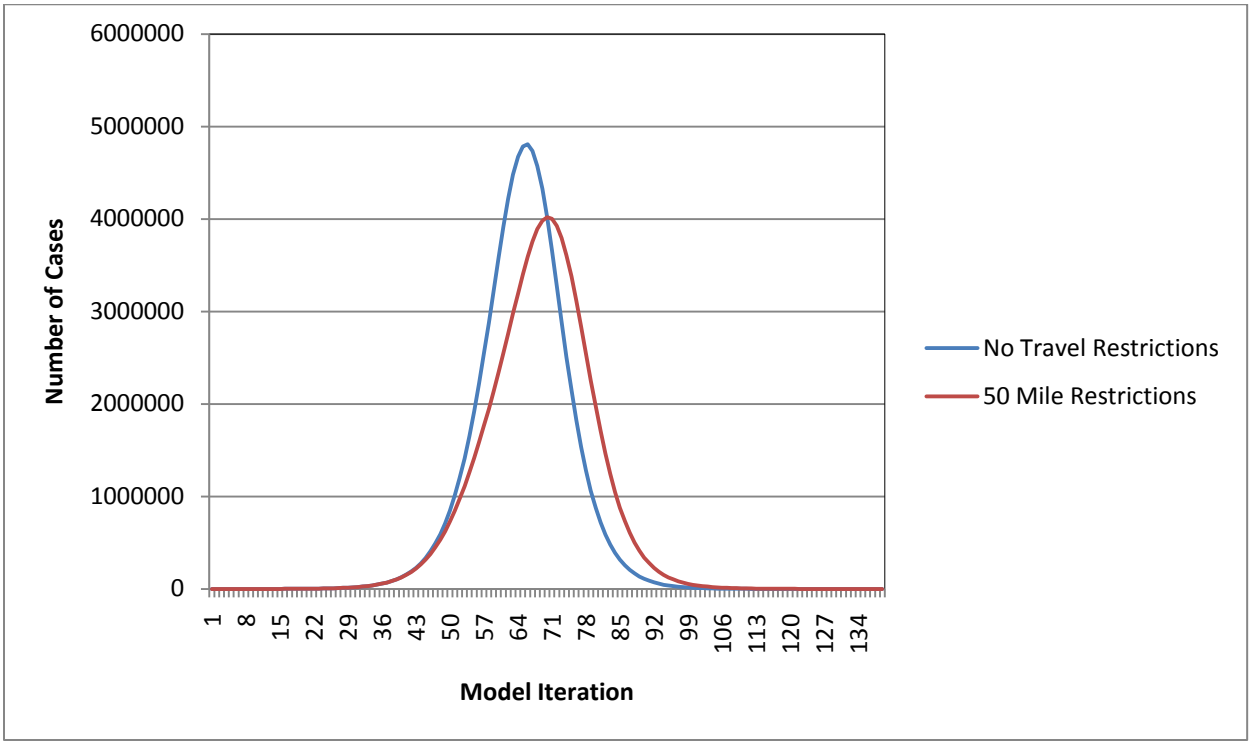

Figure 4: Impact of Travel Restrictions within the Unites States

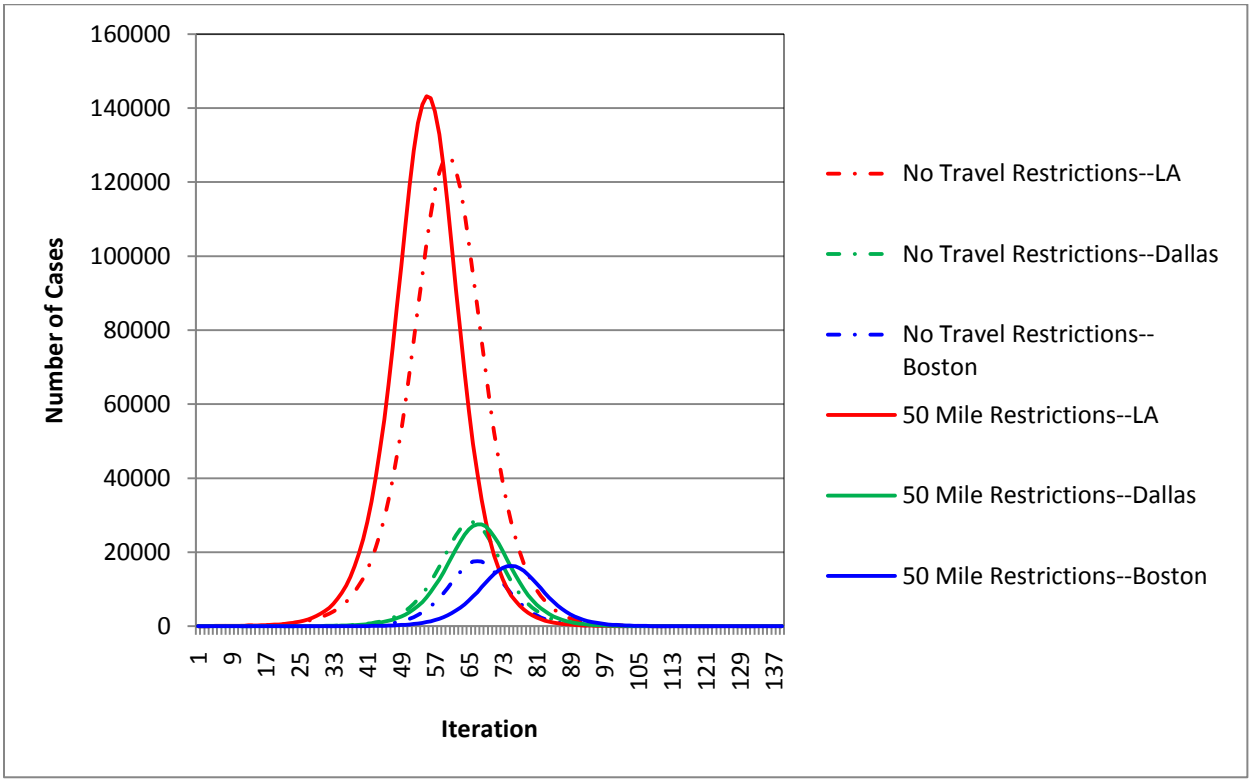

IRED Model Application: Scenario Comparison Tool

This section provides one example of how the IRED Model can be used to compare the impact of different epidemic scenarios. Consider two hypothetical epidemics: one epidemic begins with 50 initial infections all located in Los Angeles, the other epidemic begins with 15 cases in LA and another 15 cases in NYC. Which of these epidemics is worse? On the surface this looks like a simple question. However, this ignores the fact that there are many different ways to define "worst". For example, a federal public health official may be concerned with total national incidence while a Massachusetts health official may consider local cases paramount. Table 3 demonstrates that the IRED Model can compare different epidemic scenarios at different scales, a fact of great significance to local officials, who may rank epidemic outcomes differently than federal officials. On infection hierarchies, see [5]. Acknowledging these dual viewpoints will permit contingency planning that is attentive to local and national perspectives.

Table 3: Comparing Mean Number of Infections in Different Epidemics (a = 1, b = 2)

	Starting	Day 10	Day 20	Day 30	Day 40	Day 50
US	50	259	1945	14632	108923	739078
LA	50	107	436	2157	11547	54987
NYC	0	0.5	11	107	862	7786
Chicago	0	1.1	17	123	850	7218
Dallas	0	0.3	7.3	56	458	3557
Boston	0	0	1.9	13	150	1393
US	30	157	1190	8916	67513	488471
LA	15	36	143	747	4198	24010
NYC	15	12	43	203	1206	8256
Chicago	0	1 .0	18	120	753	5840
Dallas	0	0.1	3.6	38	302	2452
Boston	0	0.7	5.4	29	214	1506

Potential IRED Model Users

The IRED Model will be useful in research and applied policy settings. As a research tool, it offers a fast and efficient method for computing spatio-temproal epidemic trajectories. These, in turn, can inform policy-makers in comparing interventions and projecting impacts of a wide range of scales from local to national. This can improve the allocation of scarce resources, contributing to public health.

Conclusion

The IRED model offers a computationally efficient tool for projecting spatially heterogeneous epidemic severity from arbitrary initial conditions in a stochastic model. For every model region, confidence intervals for mean incidence can be constructed. The impact of policy interventions such as travel restrictions, or behavioral adaptations such as travel avoidance can be assessed. Epidemic severity at different scales—local and national—can be estimated. And forecasting can proceed sequentially with near-term forecasts updating as new data become available.

Acknowledgements

The authors thanks Joshua M. Epstein for insightful comments and editorial suggestions and Julia Chelen for her assistance. The project was supported by: The DHS National Center for the Study of Preparedness and Catastrophic Event Response (PACER) at the Johns Hopkins University through Award Number N00014-06-1-0991 from the Office of Naval Research; The NIH Models of Infectious Disease Agent Study (MIDAS), Award Number U01GM070708 from the National Institute of General Medical Sciences; The University of Pittsburgh Public Health

Adaptive Systems Project (PHASYS) through Cooperative Agreement Number 1P01TP000304-01 from the Centers for Disease Control and Prevention; National Science Foundation Grant Number SES-0729262 Collaborative Research: Modeling Interaction Between Individuals Behavior, Social Networks and Public Policy to Support Public Health Epidemiology. The content of this paper is solely the responsibility of the authors and does not necessarily represent the official views of any funding organization.

References

[1] Bailey, Norman T. J. (1953). The total size of a stochastic epidemic. Biometrika, 40, 177.

[2] Whittle, P. (1955). The outcome of a stochastic epidemic-a note on Bailey's paper. Biometrika, 42. 116-22.

[3] Keeling, M. J., Bjørnstad, O. & Grenfell, B. T. (2004) Metapopulation dynamics of infectious diseases. pp. 415-446 In Ecology, genetics and evolution. Burlington MA: Elsevier.

[4] Xia, Y. C., Bjørnstad, O. N. & Grenfell, B. T. (2004) Measles metapopulation dynamics: A gravity model for epidemiological coupling and dynamics. American Naturalist 164, 267-281.

[5] Grenfell, B. T. & Bolker, B. M. (1998) Cities and villages: infection hierarchies in a measles metapopulation. Ecology Letters 1, 63-70.

[6] Allen, L. J. S. and A. Burgin (2000) Comparison of deterministic and stochastic SIS and SIR models in discrete time. *Mathematical Biosciences*. 163: 1-33

[7] Grenfell, B. T. & Harwood, J. (1997) (Meta)population dynamics of infectious diseases. Trends in Ecology and Evolution 12, 395-99.

[8] Lloyd, A.L. & Sattenspiel, L. (2009). Spatiotemporal Dynamics of Measles: Synchrony and Persistence in a Disease Metapopulation. In: *Spatial Ecology*, S. Cantrell, C. Cosner and S. Ruan (eds.), Chapman and Hall.

[9] Lloyd, A.L., & Jansen, V.A.A. (2004). Spatiotemporal dynamics of epidemics: synchrony in metapopulation models. Math. Biosci. **188**, 1-16. PDF

[10] Allen, L. J. S. and A. Burgin (2000) Comparison of deterministic and stochastic SIS and SIR models in discrete time. *Mathematical Biosciences*. 163: 1-33.

[11]Allen, L. J. S. and E. J. Allen. (2003) A comparison of three different stochastic population models with regard to persistence time. *Theoretical Population Biology.* 64: 439-449.

[12] McCormack R. K. and L. J. S. Allen. 2007. Multi-patch deterministic and stochastic models for wildlife diseases. *Journal of Biological Dynamics.* 1(1): 63-85.

[13] Allen, L. J. S., Y. Lou, and A. L. Nevai (2009) Spatial patterns in a discrete-time SIS patch model. *Journal of Mathematical Biology.* 58: 339-375

[14] Kermack, W. O. and McKendrick, A. G. "A Contribution to the Mathematical Theory of Epidemics." *Proc. Roy. Soc. Lond. A* **115**, 700-721, 1927.

[15] Voorhees, Alan M., 1956, "A General Theory of Traffic Movement," 1955 Proceedings, Institute of Traffic Engineers, New Haven, Connecticut.

[16] Anderson, R. and May, R. 1993. Infections Disease of Humans: Dynamics and Control. Oxford Sciences Publications.

[17] Keeling, M. J. and Rohani, P. 2007. Modeling Infectious Diseases in Humans and Animals. Princeton University Press.

[18] Ferguson, N., Cummings, D.A.T., Fraser, C., Cajka, J.C., Cooley, P.C., Burke, D.S. 2006. Strategies for mitigating an influenza pandemic.Nature July 27;442(448-452). [PubMed: 16642006]

[19] Ferguson, N., Cummings, D.A.T., Cauchemez, S., et al. 2005. Strategies for containing an emerging influenza pandemic in Southeast Asia.Nature September 08;437(209-214). [PubMed: 16079797]

[20] Parker, J and Epstein, J.M. (2011) A Distributed Platform for Global-Scale Agent-Based Models of Disease Transmission. Forthcoming in ACM TOMACS

[21] Rvachev LA, Longini IM (1985) A mathematical model for the global spread of influenza. Math Biosci 75: 3--22

[22] Epstein JM, Goedecke DM, Yu F, Morris RJ, Wagener DK, et al. 2007 Controlling Pandemic Flu: The Value of International Air Travel Restrictions. PLoS ONE 2(5): e401. doi:10.1371/journal.pone.0000401