

On Understanding Catastrophe: The case of highly severe influenza-like illness (Position Paper)

Magnus Boman
SICS and KTH/ICT/SCS
Kista, Sweden
mab@kth.se

Abstract

Computational epidemiology is a form of spatiotemporal reasoning in which social link structures are employed, and spatially explicit models are specified and executed. We point to issues thus far addressed neither by engineers, nor scientists, in the light of a use case focusing on catastrophic scenarios that assume the emergence of a highly unlikely but lethal and contagious strain of influenza. Our conclusion is that important perspectives are missing when dealing with policy issues resulting from scenario execution and analyses in computational epidemiology.

1 Introduction

In the summer of 2010, we were asked to participate in the filming of a Swedish TV series, scheduled for nationwide broadcast and intended for subsequent international broadcasting, on the end of humanity. The threat to humanity at hand was that of communicable disease. Our roles were those of computational epidemiologists: experts on all aspects epidemiological except the medical ones. The producers knew that we had developed one of the largest simulation platforms for epidemiological studies in the world [Brouwers *et al.*, 2010], and that we had built it on registry data on the entire Swedish population, not on samples (cf., e.g., [Eubank, 2004; Ferguson *et al.*, 2005]). We agreed, on two conditions. First, the specification of the actual disease had to be made by an epidemiologist, to guarantee the plausibility of the pathogen, i.e., the morbidity and mortality had to be realistic even if the occurrence of the pathogen could have a miniscule probability. Second, our work on running simulations of more realistic scenarios— in which the contact network of all Swedes help calculate the spread of diseases with observable and measurable characteristics— should not be disturbed.

Because of the recent H1N1 pandemic, and because our platform had just been used for calculating its cost [Brouwers *et al.*, 2009], it was agreed that the pathogen should be influenza, with properties only the Spanish Flu could compare to among real pandemics. The haemagglutinin and neuraminidase combination (the H and the N of the specified influenza virus) was to be novel. Such a combination was in the

end specified, by Björn Olsen, one of the researchers that discovered H16; the latest haemagglutinin to be found [Fouchier *et al.*, 2005]. While the properties of this influenza—harrowing as they were—ultimately proved insufficient to, by themselves, cause the end of humanity, our experiments suggested to us a new line of applications for the social sciences within the realm of epidemiology. While network theory [Newman, 2003], mathematical modelling [Hufnagel *et al.*, 2004], and social mechanisms [Hedström, 2005] may help understand the processes of spread of communicable disease in society [Epstein, 1999], the sheer impact of a disease more deadly than anything experienced in history propels the results into a different scientific territory, viz. practical ethics and social psychology. In fact, our second proviso could not be met, because running these simulations provided us with some unexpected lessons and questions that affected the way we reason about computational epidemiology. The purpose of this position paper is to point to those issues, and thereby assist in hastening methodological development in this rapidly expanding area.

2 Outbreak scenario

The influenza strain specified had roughly the following characteristics (see, e.g., [Giesecke, 2002] for basics on epidemiology). On day zero, replication of the new virus sets in after a few hours, in the infected cells. When the first symptoms (sore throat or eyes, beginning headaches) appear, after 24 hours, virus levels are very high. On days 1-3 after infection, they peak and coincide with a cough. The virus levels slowly diminish, until they disappear 6-7 days later. The ability for the virus to replicate is counter-balanced by either existing antibodies or by the production of new antibodies. That process starts after a few days, leading to zero infectiousness after little more than a week, even if the cough persists. From the time of infection to the first symptoms occur, infectiousness is lower than on days 1-3, in spite of active virus replication, since the cough is absent.

The scenario was built around an airplane arriving to Arlanda airport, Stockholm, on day zero. No passenger had symptoms when boarding the plane, in South-east Asia, but one of the passengers were infected two days before boarding. On the plane, the passenger develops a cough, and infects five more passengers before the plane lands. Since the virus strain is new, there is no herd immunity to take into account. The

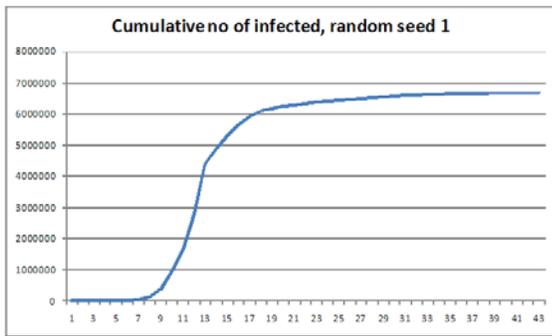


Figure 1: The number of infected in one typical run of the science fiction scenario, out of several dozens. Since this is merely for illustration, we include neither averages nor significance tests. Weeks are on the x axis and the number of infected individuals on the y axis.

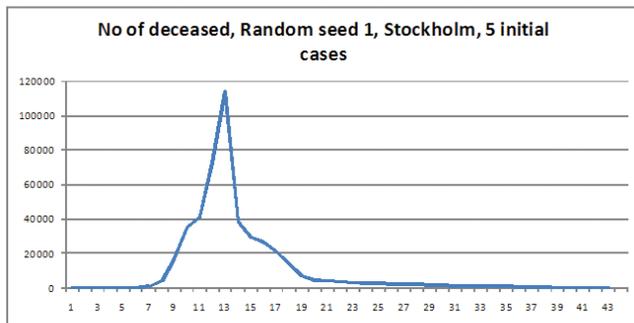


Figure 2: The number of deceased from the same run, axes as before.

attack rate is very high (around 50 per cent), and out of those infected 1-5 per cent need hospital care. Out of those with severe infections, all of which are in need of hospital care, 1-5 per cent die. In accordance with our usual workflow [Cakici and Boman, 2011], this scenario was detailed through several rounds of iteration, which provided more parameter values for input into our executable model. The latency period was 18 hours, for instance. Severely infected individuals received the highest probability of going to hospitals: 50 per cent went to primary care, 10 per cent to hospitals, 1 per cent to work, and the rest stayed home.

3 Experiments

Getting an epidemic to take off is often a problem in realistic simulation environments, since for most realistic parameter settings, there is no serious outbreak. In this case, however, the attack rate and general properties of the pathogen caused serious outbreaks in almost every run. We ran our simulations with various random seeds, affecting the stochastic parameters of the model in various ways. We also checked the stability of the results through sensitivity analyses in which

we varied an important variable representing context infections. In short, the probability of infection is highest in the family home and at work places, and lower in the neighbourhood and when travelling. These probabilities, the so-called beta values of the disease, were adapted to match the attack rate numbers provided. Then the sensitivity analyses varied the number of colleagues an individual had an opportunity to infect on a particular day, the number of people encountered in the neighbourhood, the likelihood of travel, and other stochastic parameters in our model. The reason for including results here is more pedagogical than an attempt to convey actual results, since the entire scenario is pure science fiction. That said, we ran a number of scenarios, and also varied the index cases geographically, e.g., by assuming the first six people arrived to other airports than Arlanda, and even appeared in random places in the country. In all runs, the pandemic took off when it reached Stockholm, where approximately one sixth of Sweden’s population of about nine million resides.

4 Results

The number of infected in our simulations were about 6.5 million, corresponding to approximately 73 per cent of the Swedish population of 8.861.388. (The current population is over nine million, but our population data is not quite up to date, hence the difference.) Out of those, 7.978.105 individuals belong to the giant component, i.e. just over 90 per cent. The number of infected is shown in Figure 1 in cumulative form. The graph representing the deceased is extremely peaked, see Figure 2. In week 13 in this particular run, almost 115.000 people died, and almost 1.600.000 were infected. The sharp decline after the peak is to a large extent due to the network properties: Sweden is a sparsely populated country and of the 90 per cent part of the giant component, very few had not been exposed by then. Already in week 13, half the entire population were infected. By week 18, the number of infected passed 6.1 million. When we ended the simulation run after week 43, less than 1500 new cases occurred each week, chiefly due to the fact that the population had been exhausted.

5 Discussion and conclusion

First, we must note that in this short paper, we have left out almost all details of how our model is specified and implemented. This has been published elsewhere, and the platform was open-sourced in 2010 under a GPL license [SMI, 2010]. The scientific track record of our platform is, however, not pivotal to the points we want to make. Instead, we would like to focus on the fact that we thought that the cross-disciplinary team we have been a part of for the last seven years or so—constructing and implementing the software, as well as applying it to a number of policy problems—sufficed for interpreting the output of pretty much any simulation run, in any scenario. This time, though, it was different.

Upon completion of our sets of runs, and the usual post-processing of output (graphing the distributions and looking at the logs, etc.), we presented our results to the epidemiologist and the producers of the TV series. We discussed the

outcome, and in the process also discussed the same technical issues as for an ordinary scenario: the geographical spread of the disease, the effects of travel, possible policy measures (shutting down schools, reducing operations at airports and train stations, vaccination), and the limitations (no reservoirs modelled, for instance). Quickly, however, the discussion turned to different matters, which were more of a societal nature.

We had previously been aware of cultural differences between Sweden and other countries in the face of pandemics. Naturally, the computational epidemiologist would like to realistically model all forms of social interaction (cf. [Boman *et al.*, 2006]). In the case considered here, potentially awkward cultural and societal questions surfaced, such as how and where to bury the dead. Just like the number of people capable of distributing vaccine—and the geographical location of vaccine storage facilities—constitute logistic constraints on efficient vaccine distribution, the procedures for handling the deceased may constrain efforts to diminish the effects of the disease. The number of infected is one thing, but what about the number of *affected*? With so many families affected by ill health, and even death, what chance does a public health infrastructure have to withstand the pressure of the disease? Even a resilient society could face a phase transition of sorts, in which important societal functions ceased to operate. And even if this is not the case, even if society withstands the horror of this kind of virus and quickly bounces back, as after the Spanish flu (and the added horror of WWI in that case), which is the proper scientific forum for this kind of discussion? And how does the results of such discussions propagate back into computational epidemiology and into policy measures, and ultimately politics? How does one prevent unnecessary concern, fear, or even panic?

The short answer is perhaps ethics, and practical ethics in particular. A broader perspective would include social psychology and philosophy (decision theory and value theory being included). These perspectives are then added to those that already are covered by the areas of epidemiology, computer science, sociology, statistics, mathematics, economics, and political science; all of which constitute parts in most of the largest modelling and policy projects currently in existence (see, e.g., [Ferguson *et al.*, 2006; Hall *et al.*, 2007; Halloran *et al.*, 2008]).

Our conclusion must be that important perspectives are missing when dealing with policy issues resulting from scenario execution and analyses in computational epidemiology. With many basic scientific questions already pressing (e.g., verification and validation, replication and docking, and security and privacy issues), this might be hard to fit on the agenda of pandemic preparedness. We would like to open up for discussion now, before anything even remotely like what we have just simulated is encountered by a human population.

Acknowledgements

The ASSETS European Union Best Practice Network helped finance the link structure aspects of this research, and we thank our colleagues in S-GEM [SGE, 2011]. In particular, Baki Cakici provided important technical assistance.

References

- [Boman *et al.*, 2006] M. Boman, A. Ghaffar, F. Liljeros, and M. Stenhem. Social network visualization as a contact tracing tool. In N. et al. Jennings, editor, *Proc WS on Agent Technology for Disaster Management*, pages 131–133. Hakodate, Japan, 2006.
- [Brouwers *et al.*, 2009] L. Brouwers, B. Cakici, M. Camitz, A. Tegnell, and M. Boman. Economic consequences to society of pandemic H1N1 influenza 2009: Preliminary results for Sweden. *Eurosurveillance*, 14(37), 2009.
- [Brouwers *et al.*, 2010] L. Brouwers, M. Boman, M. Camitz, K. Mäkilä, and A. Tegnell. Micro-simulation of a smallpox outbreak using official register data. *Eurosurveillance*, 15(35), 2010.
- [Cakici and Boman, 2011] B. Cakici and M. Boman. A workflow for software development within computational epidemiology. *J of Computational Science*, in press, 2011.
- [Epstein, 1999] J. M. Epstein. Agent-based computational models and generative social science. *Complexity*, 4(5):41–60, 1999.
- [Eubank, 2004] S. Eubank, *et al.* Modelling disease outbreaks in realistic urban social networks. *Nature*, 429:180–184, 2004.
- [Ferguson *et al.*, 2005] N. M. Ferguson *et al.* Strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature*, 437(7056):209–214, 2005.
- [Ferguson *et al.*, 2006] N. M. Ferguson *et al.* Strategies for mitigating an influenza pandemic. *Nature*, 442:448–452, 2006.
- [Fouchier *et al.*, 2005] R. A. M. Fouchier *et al.* Characterization of a novel influenza a virus hemagglutinin subtype (H16) obtained from black-headed gulls. *J. Virol.*, 79(5):2814–2822, 2005.
- [Giesecke, 2002] J. Giesecke. *Modern Infectious Disease Epidemiology*. Arnold Publishers, 2nd edition, 2002.
- [Hall *et al.*, 2007] M Hall, R Gani, H E Hughes, and S Leach. Real-time epidemic forecasting for pandemic influenza. *Epid Inf*, 135(3):372–385, 2007.
- [Halloran *et al.*, 2008] M. E. Halloran *et al.* Modeling targeted layered containment of an influenza pandemic in the united states. *PNAS*, 105(12):4639–4644, 2008.
- [Hedström, 2005] Peter Hedström. *Dissecting the Social*. Cambridge Univ Press, 2005.
- [Hufnagel *et al.*, 2004] L. Hufnagel, D. Brockmann, and T. Geisel. Forecast and control of epidemics in a globalized world. *PNAS*, 101:15124–15129, 2004.
- [Newman, 2003] Mark E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.
- [SGE, 2011] S-GEM: Stockholm Group for Epidemic Modeling, 2011. <http://s-gem.se>.
- [SMI, 2010] MicroSim source code, October 2010. <https://smisvn.smi.se/sim/>.