

# **Agent-Based Models As Policy Decision Tools: The Case of Smallpox Vaccination**

Till Grüne-Yanoff

Royal Institute of Technology, Stockholm

## **1. Introduction**

Imagine that smallpox viruses were intentionally released in an industrialized country. What should the government do to mitigate the effects of such an attack? With the recent surge in security concerns, this question has gained some urgency. The answer is not straightforward. The only known medication against smallpox is vaccination with cowpox virus. But a significant minority of individuals experience severe (and often fatal) adverse effects to the vaccine, such that a national preventive mass vaccination may cause more deaths than an isolated smallpox epidemic (Kemper et al. 2002). To make a rational policy decision, more information is needed about the potential development of an epidemic, and about the effects of different vaccination policies.

Unfortunately, the available empirical data is limited. Smallpox outbreaks have not occurred in the western world since 1972; virus is eradicated since 1977. Most of WHO's policies, which led to its eradication, were employed in societies very different from those of the industrialised world. To make up for this deficit, scientists have built computer models that simulate smallpox outbreaks in relevant contexts. In particular, a number of scientists over the last five years have employed agent-based models (ABM) to simulate smallpox epidemics and assess policy options (Halloran et al. 2002, Epstein et al. 2004, Eidelson and Lustick 2004, Eubank et al. 2004, Burke et al 2006, Brouwers 2006). These models explicitly track the progression of the disease through each individual, and track the contacts of each individual with others in the relevant social networks and geographical areas. Epidemiological models thus acquire a social science component. It has been argued that this component is central for making rational policy decisions, and indeed, the policy recommendations based on ABM differ considerably from models that do not take a social component into account.

This paper investigates how ABM support smallpox policy decisions, and what sort of decision-making procedures it supports. It uses the smallpox studies as case studies of such ABM supported decisions. Section 2 discusses the general ABM approach to smallpox.

Section 3 surveys five agent-based smallpox models. Section 4 discusses four decision procedures and their applicability in making policy choices from ABM. Section 5 concludes.

## 2. Modelling epidemics

ABM are used as decision tools for a wide variety of real-world problems, for example crowd control, minority segregation or flood management. I will focus on the smallpox case for the following reasons. First, it offers clear and well-understood decision alternatives. Second, the objectives are unambiguous: what counts is saving lives at reasonable costs. Third, a number of studies with similar focus but differing models have recently been published. Their similarities and differences illustrate well the specific issues with ABM. Last, authors of these papers explicitly discuss the social assumptions of their models.

The smallpox policy interventions most prominently discussed are mass vaccination and trace vaccination. A mass vaccination (MV) involves vaccinating a large part of the population, either before any cases of smallpox are confirmed (preventive MV) or after the first confirmed case (post-release MV). In trace vaccination (TV), also called ring vaccination or targeted vaccination, every contact of a confirmed smallpox case is traced and vaccinated. The CDC defines contacts as ‘Persons who had ... close-proximity contact (<2 meters) with a confirmed or suspected smallpox patient after the patient developed fever and until all scabs have separated’ (CDC 2002).

These policy options were first discussed on the basis of standard epidemiological models. An often-cited example is Kaplan et al. (2002), which simulates an attack of 1000 initial smallpox cases on a population of 10 Million. The population is assumed to mix homogeneously – i.e. to consist of identical individuals, who have an equal chance of interacting with any other population member.  $R_0$ , the rate of infections a single infectious agent generates among susceptibles, is assumed to be uniform throughout the simulation.  $R_0 = 3$  is derived from historical data. An infected agent undergoes four stages. Only in the first is she vaccine-sensitive; only in the third and fourth is she infectious; in the fourth, however, she shows symptoms (scabs) and is automatically isolated. Additionally, administration of vaccination is modelled under logistical constraints: MV of the whole population is achieved in ten days. Tracking *and* vaccinating an infected person in TV, however, takes four times as many nurse-hours as just vaccination. The model is deterministic (it does not include any random variables). The results show a wide difference for simulations run with TV and MV,

respectively. Both initiated on day five after the initial attack, MV leads to 560 deaths, while TV leads to 110,000 deaths. Sensitivity analysis shows that TV is more sensitive than MV to size of initial attack and changes in  $R_0$ , further supporting the strong results in favour of MV.

Models of this kind been criticised for neglecting important causal mechanisms of an epidemic. Smallpox is spread almost exclusively by extended face-to-face contact. Therefore, it is not just the number of infected individuals that matter, but the patterns of exposure. These patterns are constituted by the social interactions of infectious and susceptible individuals. Explicitly modelling social contacts can have crucial impact on the result of a model:

‘changing the pattern of connections between exposed and unexposed individuals can often affect population infection levels more than changing the exposure status of individuals in that population’ (Koopman and Lynch 1999, 1170)

The so-called small-world effects nicely illustrate the importance of such patterns. To model them, social connections are represented as a network on a population. Assuming that each population member has the same number  $2K$  of contacts, these contacts can be arranged in different ways. A maximally ordered network consists of a ring, on which each agent is connected to her immediate  $K$  neighbours to her left and right. Less ordered networks are achieved by ‘rewiring’ each link of every agent (in clockwise order) with probability  $p$  to any other agent in the system. Thus a maximally ordered network is characterised by  $p = 0$ , and a complete disordered network (similar to the case of homogeneous mixing) is characterised by  $p = 1$ . Kuperman and Abramson (2001) simulated the transmission of a (non-fatal) epidemic with a number of stages in such networks. They found that the evolution of epidemics undergoes a phase transition at some degree of disorder: while networks with small  $p$  produce a low-amplitude evolution, networks with large  $p$  produce wide and very regular oscillations. In such abstract investigations, the kind of social connection occurring in a population greatly matters.

To include social factors into concrete epidemic models is easier said than done. It requires substituting a black box with an explicit decision-making model for each agent. Spaces in which agents move have to be accounted for, as well as the contacts and connections agents make in these spaces. Given a lack of both data and theory appropriate for such an undertaking, such modelling projects are beset with severe problems.

Nevertheless, several smallpox ABM have been developed. They build the epidemic macro-phenomenon ‘from the bottom up’ - by designing agents with heterogeneous attributes (age, identity) and heterogeneous behavioural rules (conditional on individual attributes) interacting in discrete time in an explicitly modelled landscape. Epidemic spreads develop from these modelled interactions, taking the infectious properties of particular places, kinds of interactions and kinds of agents into account. Five such models are surveyed in the next section.

### **3. Five case studies of smallpox models**

(i) Halloran et al (2002) simulate an attack of 10 infected agents on a stochastically (through random draws) generated community of 2000.<sup>1</sup> The community consists of four neighbourhoods, each containing large day-care centres and small playgroups. Two neighbourhoods share an elementary school, and all four share a middle and a high school. Households in each neighbourhood contain up to seven people. Agents are characterised by age, family ID, disease status and vaccination status. Household size and age distributions are based on the US 2000 census. During the course of a ‘day’, agents visit different spaces according to age and family ID, and then return home. If sick, agents stay home with a probability conditional on their age. Within the visited spaces, they meet others with a fixed probability. The model thus is a variant of the homogeneous mixing hypothesis, only separated into smaller subgroups.

Transmission rates from infectious to susceptible agents are dependent on place of contact and age of the agents involved. Halloran et al. first derive a general household transmission rate from African/Asian data of the 60s and 70s. On that basis, they then derive relative magnitudes for adult-adult, adult-child, etc. transmissions from influenza data. Transmission rates for other places of contact were derived from influenza infection rates.

Because of its stochastic nature, the results from the simulated intervention can only be compared with respect to their distributions, not single numbers. 200 simulations were performed for each intervention. At 80% MV after the first confirmed case, on average 0.9

---

<sup>1</sup> An extension of the model for larger populations and with a more detailed disease natural history is presented in Longini et al. (2006)

out of 1000 individuals died, while at 80% TV, 10.9 out of 1000 died. Similar differences were found if vaccinations were administered after the 15<sup>th</sup> (9.4 vs. 19.6 deaths) and the 25<sup>th</sup> case (13.7 vs. 28.2 deaths). However, while MV prevented 0.5 cases per dose given, TV prevented 2.01 (both administered after 1<sup>st</sup> case).

(ii) Burke et al (2006), based on Epstein et al (2004), simulate a single initial infected person attack on a town network of either 6.000 or 50.000 people. Town networks either consist of one town (uniform), a ring of six towns, or a ‘hub’ with four ‘spokes’. Each town consists of households up to seven persons, one workplace and one school. All towns share a hospital. Each space is represented as a grid, so that each cell in the grid has eight neighbours (so-called Moore neighbourhood). Agents are distinguished by type (child, health care worker (5% of adult population), commuter (10%) and non-commuter (90%)), by family ID and by infectious status. Each ‘day’, agents visit spaces according to their type, and then return home. On the first ‘day’ of the simulation, the position in schools and workplaces is randomly assigned, but after that agents remember their positions. During a ‘day’, agents interact with all of their immediate neighbours: 10 times at home, 7 times at work and 15 times in hospital. After each interaction, they will move positions to the first free cell in their neighbourhood. Homogeneous mixing is thus completely eschewed; instead, agents interact in a number of dynamic neighbourhoods.

Transmission occurs with a certain rate at each of the agents’ interactions. It can both infect contactor and contactee. Transmission rates depend on the stage the infectious person is in, the type of disease he has, and whether the susceptible agent has partial immunity.

Burke et al. assess only TV as a first policy intervention, and MVs of varying degrees only as ‘add-on’ measures. Results for all three town networks showed substantial concordance. Contrasted with a ‘no response’ scenario, TV in combination with hospital isolation was sufficient to limit the epidemic to a mean of fewer than 48 cases and a mean duration of less than 77 days. Post-release MV of either 40% or 80% of the total population added some additional protection, reducing the mean of infected people to 33 and shortening the mean duration to less than 60 days.

(iii) Identity Neighbourhoods

Eidelson and Lustick (2004) simulate a single infected agent attack on a population of 1764. Their model consists of a 42×42 cell lattice, representing social relationships, not geographic ones. Agents occupy a fixed position on the lattice. An agent is characterised by its coordinate address, its identity repertoire (a finite set of affiliations, represented by integers) and its currently activated identity. At each time step, identities change in response to neighbours' identities. If sufficiently many neighbours display an identity that is in the agent's repertoire, but is not her current identity, the agent will switch to this identity. If it is not in her repertoire but above a certain threshold, she will adopt the identity into her repertoire. If it is above a higher threshold, she will also switch to this identity.

Transmissions occur between direct neighbours. Transmission probabilities depend on the stage an infected agent is in, but particularly on whether a susceptible agent shares identities with an infected agent. Once infected, the agent goes through the various disease stages, becoming increasingly infectious to others, but also becoming increasingly likely to be detected and to be removed from the population.

Simulation results show that TV rivalled MV if a very substantial proportion of smallpox cases could be detected and isolated almost immediately after infection, or if residual herd immunity in the population was relatively high. At 80% MV, 15 days after the first case, ca. 3.5% of the population is infected, while at 80% TV, ca. 6% is infected.

(iv) Eubank et al (2004) simulate an attack of 1000 infected agents on the population of Portland, OR, of 1.5 Million. Portland is represented by ca. 181,000 locations, each associated with a specific activity, like work, shopping, school, etc. and maximal occupancies. Each agent is characterised by a list of the entrance and exit times into and from a location for all locations that person visited during the day. This huge database was developed by the traffic simulation tool TRANSIMS, which in turn is based on US census data.

Smallpox is modelled by a single parameter, disease 'load' (analogous to a viral titre). Agents have individual thresholds above which their load leads to infection (and load growth at individual growth rates), symptoms, infectiousness and death. Infectious agents shed a fixed fraction of their load to the local environment per hour. Locations thus get contaminated with load, which is distributed equally among those present. Shedding and absorption fractions

differ individually. Infected individuals withdraw to their homes 24h after becoming infectious.

Eubank et al (2004) model is deterministic. MV with a 4-day delay resulted in 0.39 deaths per initially infected person; TV with the same delay in 0.54 deaths. Varying delays, they found that delay in response is the most important factor in limiting deaths, yielding similar results for TV and MV.

(v) Brouwers et al. (2006) simulate an attack of 50 infected individuals on the population of Stockholm plus 1/10<sup>th</sup> of the Swedish population.<sup>2</sup> The model is based on a 100×100 meter grid of Sweden, which contain households, kindergartens, schools, offices and hospitals according to the 1998 census. Agents are characterised by their age, sex, family-ID, workplace, immunity, disease stage and hospital affiliation.

Transmission from infected to susceptible agents in a place of contact depends on a predefined transmission probability for that type of place, the stage in which the infectious person is in, and the number of infected present at this place. For large places, the number of maximal contacts is restricted.

Brouwers et al. simulate four policy interventions: Hospital staff vaccination, TV, MV and a combination of TV and hospital staff vaccination. While TV yields a considerably higher number of infections than MV (average 34,77 vs. 17,43) the authors argue that this difference is insignificant in the light of the average of needed vaccinations (125 vs. 6897378), and that therefore TV is preferable to MV.

#### **4. Making decisions with ABM**

The surveyed papers give more or less unconditional policy advice. For example: ‘contact tracing and vaccination of household, workplace and school contacts, along with effective isolation of diagnosed cases, can control epidemics of smallpox’ (Burke et al. 2006, 1148); ‘outbreaks can be contained by a strategy of targeted vaccination combined with early detection without resorting to mass vaccination of a population’ (Eubank et al. 2004, 180).

---

<sup>2</sup> The authors plan future simulations for the whole population of Sweden.

The measures recommended in these quotes evidently refer to measures taken by real policy makers, not to interventions in a model. Yet the recommendations are based only on the results that simulated interventions had in simulated worlds. So how can a (real) policy maker justify a policy choice on the basis of such simulation results? I discuss four decision procedures and their appropriateness for this question.

The simplest justification is to claim that the model represents all the relevant causal relationships of the world, and thus can be treated as a surrogate for the real system. Eubank et al. (2002) may be interpreted this way: it simulates the outbreak in the minutely represented city space of Portland, involving city inhabitants whose every move is determined by data derived from recent censuses, and it models the development deterministically (i.e. without any random variables). If the model were correct, a policy maker could take the predictions the model makes for each policy intervention, rank the outcomes, and choose the policy with the best outcomes.

However, the model contains many unwarranted assumptions. For example, the authors make uniform assumptions about the occupancy rate of locations within a city block that are, the authors admit, ‘nothing more than reasonable guesses’ (Eubank et al 2004, supplement). Location occupancy rates crucially influence the number of possible contacts (and hence the disease spread). If such a parameter is determined by guess, then there is no reason to believe that the presented model gives accurate point predictions, or can function as a surrogate for the real system. Making a policy decision on the basis of such a model alone is not rational.

By building a deterministic model, Eubank et al. (2004) were forced to face head-on the lack of detailed data about specific locations and individuals. Often, all that is known (from census and land use data) are the distributions of such attributes across all agents and locations. Here, stochastic modelling is appropriate. A stochastic model draws, for each simulation, the specific attribute values for each individual and space from the known distributions. In effect, this means that each simulation has its own model: similar in structure, but with different parameter realisations. All papers surveyed here except for Eubank et al. (2004) use stochastic modelling. They create a large number of model realisations (from 35 to 200), simulate different vaccination policies on them, and obtain a distribution of outcomes for each policy. These distributions are commonly only described by their average or mean (often not clarified



which) or by their mean and standard deviation (only in Eidelson and Lustick 2004 and Burke et al. 2006).

If the distributions could be quantified, so that each model specification had a determined probability assigned, then the standard decision procedure under risk, expected utility maximisation, could be applied. The policymaker could list for each model the outcomes of each vaccination policy; she could then rank these outcomes with the help of a utility function (as a proxy, number of deaths may serve as such a numerical ranking). Last, she could sum up the utilities of the outcomes of each policy, weighing them with the probability of the model in which they occur. This would yield a single expected utility number for each policy, according to which the policymaker could decide.

Given that the distributions of model attribute values are known, it may seem plausible that the probability of models is known. This, however, would neglect important features of models that are not constructed stochastically, in particular the model's structure and the behavioural rules of the agent. The model structure specifies which attributes are included and which not. Because there is no well-confirmed general theory of social interactions and connections, quantifying the probability with which a certain model structure is true is out of reach of even the most optimistic model builders. The behavioural rules specify agents' behaviour in the model. Again, well-confirmed general theories on this topic do not exist. In addition, survey data is not available or cannot be operationalised. Thus, one cannot quantify distributions over various behavioural rules either. Both of these features, however, are central to smallpox ABM. Assigning probabilities to stochastic models is therefore not possible.

If the parties to the decision do not know the model structure or the prior probabilities of the model parameters, situations of 'deep uncertainty' arise (Lempert 2002, 7195). Under deep uncertainty, predicting the outcomes of vaccination policies with certainty or predicting their probabilities is impossible. The expected utility method of decision-making therefore cannot be applied to them.

Some of the authors seem to acknowledge this limitation, while insisting that their models still contribute to policy decision-making:

‘We caution that the numbers of cases generated in various scenarios should not be taken as quantitative predictions, but instead as a basis for comparing and evaluating different intervention strategies.’ (Burke et al. 2006, 1148)

How is this feasible? Under deep uncertainty, models of uncertain standing produce outcomes with uncertain relevance. Instead of predicting *the* future of the system with one model, investigations reveal a ‘landscape of plausible futures’ (Bankes, Lempert and Popper 2001, 73) from a set of models deemed acceptable. For example, ‘we have designed a model that ... allows us to vary poorly understood properties through reasonable ranges’ (Eubank et al. 2004, 183). Such a modelling strategy has been described as *exploratory modelling*.

‘Modellers ... make guesses at details and mechanisms...[perform] a computational experiment that reveals how the world would behave if the various guesses were correct... [and use] series of such computational experiments to explore implications of varying assumptions and hypotheses’ (Bankes 1993, 435)

All the surveyed studies make such guesses. But while they highlight the use of empirical data, in particular the modelling of ‘real’ systems like Portland or Sweden, they say very little about the way they arrive at the assumptions that are not supported by empirical data. In particular, the behavioural rules are often not well motivated. The more methodology conscious authors argue that ‘we include the level of social detail that we believe necessary to capture the transmission dynamics of smallpox’ (Burke et al. 2006, 1148), but they do not discuss the basis for their intuitions.

Even though all studies make such guesses, an important difference lies in how the guesses are generated. Most authors seem content to present a credible world scenario in which agents’ behavioural rules exhibit some resemblance to the readers’ own lives. Only Eidelson and Lustick (2004, 2.4-2.6) derive the crucial behavioural assumptions in their model from a general theory. This way, exploratory modelling gains a degree of systematicity that the other studies lack. Eidelson and Lustick’s simulation unpacks the implications of their theoretical hypotheses. If implications are found untenable (for example, if verification fails), the authors can go back to the theory, which provides constraints on how alternative hypotheses can be constructed. In contrast, if authors change their credible world scenarios, they are making *ad*

*hoc* adjustments. Using a theory, in order to give systematicity and constraint to the generation of model assumptions, is therefore desirable in exploratory modelling.

The result of exploratory modelling is a set of plausible futures. How can the policy maker base her decisions on such a set? Two different strategies have been discussed. The first focuses on worst-case scenarios, against which policies should be hedged. Such a decision procedure ‘develops an assortment of plausible worst cases [which] can be very useful for designing hedging strategies’ (Bankes 1993, 440). Each plausible model assigns an outcome to each vaccination policy. Instead of comparing the policies in all outcomes, the policymaker may ‘hedge’ and only compare the policies with respect to their worst outcomes. The policy maker thus applies the *maximin* procedure: she chooses that policy that maximises the minimal (worst) outcome.

The problem with *maximin* is its extreme sensitivity to the number of plausible futures considered. The wider the scope, the more likely the inclusion of some outlandish terrible future. Given the uncertain status of many model specifications, exploratory modelling is prone to misspecifications. Such a misspecified model may provide the worst outcome, and the sole basis for the decision. Choosing policy with *maximin* is highly vulnerable to such misspecifications and requires great care in the selection of plausible models.

An alternative decision procedure pays equal attention to all models, and chooses that policy which performs relatively well, compared with the alternatives, across the range of plausible futures. This robustness analysis of policies (Lempert 2002, 7196) can be interpreted in two ways. If ‘performs relatively well’ means that a certain policy *P* performs better than other policies for *each* plausible future – i.e. if the other strategies are dominated by *P* – such a decision procedure is a special case of the common expected utility framework. However, it is highly unlikely that there will be only one non-dominated policy, when comparing policy outcomes in many possible futures (in the surveyed articles, the number lies between 35 and 200). Instead, ‘performs relatively well’ must mean that the policy is *satisficing*, i.e. its results, for each plausible future, lie above a certain minimal satisfaction threshold. Inevitably, such a decision procedure gives equal weight to every considered future. This may be unwarranted, but cannot be avoided because additional information is lacking. From all the procedures discussed so far, this one seems to me the most fitting for decisions based on

ABM. Great care has to be put, however, on finding all relevant futures and generating results for each of them.

So far, all envisioned policy decisions consisted of pure apriori reasoning. The models were constructed from prior knowledge that defined what was plausible/possible. It was put into a new shape (agents' attributes and behavioural rules, disease parameters) and confronted with other knowledge parts (spatial structure of the environment) to transform it into a more useful form (plausible futures of a smallpox attack in that environment). The form and information content of these plausible futures, as I argued, limited the use of standard decision procedures severely.

But maybe, apriori reasoned decisions and policy rankings are not so important, anyway. Should an epidemic occur, there will be plenty of evidence which subset of futures it is going to realise. Equipped with computing power and sophisticated statistics, policy makers will be able to 'swim ahead of the wave', predicting the immediate development of the disease with sufficient accuracy to make their policy decisions. To do so, policy makers need to be informationally prepared. ABM, instead of deciding policy choices apriori, may be of great help in such preparatory planning. They allow the systematic collection of plausible scenarios and provide the conceptual framework to represent all potentially useful available information long before an outbreak. Often, more information can be captured in a set of alternative plausible models than can be captured by any individual model. If an outbreak occurs, it can be swiftly determined which of the many parameters and release scenarios explored in the preparatory planning is actually realised, and thus which set of policies is likely to be optimal. Under this interpretation, the importance of the surveyed studies lies not in their policy recommendation, but in providing alternative models and model platforms on which current information storage can be maximised. The decisions the ABM support will be taken later, once the outbreak has occurred and the epidemic is developing.

## **5. Conclusion**

Drawing on a survey of five smallpox ABM, I discussed which decision procedures a policy maker could rationally apply to exploit the information provided by the simulations. Given the input used in the surveyed studies, I argued that expected utility maximisation and maximin procedures cannot be reasonably applied, and that robustness analysis must be

treated with caution. I further suggested that the purpose of the surveyed ABM may not consist in deciding between policies now; but rather that the ABM are good ways of representing current imperfect information for the case of a smallpox outbreak. While the discussion centred on smallpox models, its conclusions may contribute to a better understanding of the function of ABM in policy decisions more generally.

## 6. References

Bankes, Steve (1993) 'Exploratory Modelling for Policy Analysis', *Operations Research* 41(3): 435-449.

Bankes, Steven C., Robert J. Lempert, and Steven W. Popper (2001) 'Computer-Assisted Reasoning', *Computing in Science and Engineering*: 71–77.

Brouwers, Lisa, Kalle Mäkilä and Martin Camitz (2006) 'Spridning av smittkoppor – simuleringsexperiment', *SMI-Rapport Nr 5: 2006 T*

Burke DS, Epstein JM, Cummings DA, Parker JI, Cline KC, Singa RM, Chakravarty S. (2006) 'Individual-Based Computational Modeling Of Smallpox Epidemic Control Strategies', *Academic Emergency Medicine* 13(11): 1142-9, plus technical appendix at <http://www.aemj.org/cgi/data/j.aem.2006.07.017/DC1/1>

CDC 2002, *Smallpox Response Plan and Guidelines* (Version 3.0), Guide B, Part 2, <http://www.bt.cdc.gov/agent/smallpox/response-plan/files/guide-b-part2of3.pdf>, last accessed 27.02.2007

Eidelson, Benjamin M. and Ian Lustick (2004) 'VIR-POX: An Agent-Based Analysis of Smallpox. Preparedness and Response Policy', *JASSS* 7(3).

Epstein JM, Cummings DAT, Chakravarthy S, Singa RM, Burke DS (2004) *Toward A Containment Strategy For Smallpox Bioterror: An Individual Based Computational Approach*. Brookings Monographs.

Eubank, S., Guclu, H., Kumar, V. S. A., Marathe, M., Srinivasan, A., Toroczcai, Z. & Wang, N. (2004). 'Modelling Disease Outbreaks In Realistic Urban Social Networks', *Nature* 429: 180-184, plus supplement at <http://www.nature.com/nature/journal/v429/n6988/extref/nature02541-s1.htm>

Kaplan, Edward H., David L. Craft, and Lawrence M. Wein (2002) 'Emergency response to a smallpox attack: The case for mass vaccination', *Proceedings of the National Academy of Sciences* 99(16): 10935-10940.

Kemper, A. R., Davis, M. M. & Freed, G. L. 'Expected Adverse Events In A Mass Smallpox Vaccination Campaign'. *Effective Clinical Practice*. **5**, 84–90 (2002)

Koopman J, Lynch J, 1999, 'Individual Causal Models And Population System Models In Epidemiology', *American Journal of Public Health* 89: 1170 - 1174

Kuperman, M. and G. Abramson (2001) 'Small World Effect in an Epidemiological Model', *Physical Review Letters*, 86(13): 2909-2912.

Lempert, Robert. J. (2002) 'New Decision Sciences for Complex Systems', *Proceedings of the National Academy of Sciences* 99, supplement 3: 7309–7313.

Longini IM, Halloran ME, Nizam A, Yang Y, Xu S, Burke DS, Cummings DAT, Epstein JM (2007) 'Containing a large bioterrorist smallpox attack: a computer simulation approach', *International Journal of Infectious Diseases* 11: 98-108.