

Simulation & Gaming

<http://sag.sagepub.com/>

Agent-Based Models as Policy Decision Tools: The Case of Smallpox Vaccination

Simulation Gaming 2011 42: 225 originally published online 7 September 2010
DOI: 10.1177/1046878110377484

The online version of this article can be found at:
<http://sag.sagepub.com/content/42/2/225>

Published by:



<http://www.sagepublications.com>

On behalf of:

Association for Business Simulation & Experiential Learning
International Simulation & Gaming Association
Japan Association of Simulation & Gaming



North American Simulation & Gaming Association
Society for Intercultural Education, Training, & Research

Additional services and information for *Simulation & Gaming* can be found at:

Email Alerts: <http://sag.sagepub.com/cgi/alerts>

Subscriptions: <http://sag.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://sag.sagepub.com/content/42/2/225.refs.html>

Agent-Based Models as Policy Decision Tools: The Case of Smallpox Vaccination

Simulation & Gaming
42(2) 225–242
© 2011 SAGE Publications
Reprints and permission: <http://www.sagepub.com/journalsPermissions.nav>
DOI: 10.1177/1046878110377484
<http://sg.sagepub.com>



Till Grüne-Yanoff¹

Abstract

Agent-based simulation (ABS) studies have recently been employed to support policy decisions. This article addresses the particular potentials and problems that ABS faces in this usage. First, the author warns against taking “familiarity” with specific ABS as a criterion for having confidence in the model’s policy recommendations. Second, he shows that specific epistemic issues—in particular the high number of detailed simulated systems—require additional reflection on which decision rules to choose for policy decisions based on ABS. Third, the author points out directions in which the construction and uses of ABS in policy decision could be improved. Each of these issues is illustrated by simulation studies undertaken to investigate smallpox vaccination policies.

Keywords

agent-based simulations, decision making, epidemics, exploratory modeling, methodology, smallpox, vaccination policy, validity

As part of the broader trend toward evidence-based policy making, scientists have begun designing agent-based simulation (ABS) studies with the explicit goal of supporting policy decision making. ABS has recently been used to develop vaccination policies, crowd control measures, evacuation plans, market designs, industrial and fiscal policies, and many more. This contrasts with the common practice of using highly stylized, analytically tractable models, and introduces new epistemic problems, which I discuss in this article.

¹University of Helsinki, Helsinki, Finland

Corresponding Author:

Till Grüne-Yanoff, Helsinki Collegium of Advanced Study, Fabianinkatu 24 (P.O. Box 4),
00014 University of Helsinki, Helsinki, Finland
Email: till.grune@helsinki.fi

Unlike standard models, ABS simulation studies often incorporate vast amounts of detail, allow modeling and computing levels of complexity unreachable for unaided human capacities, and present their results in easily comprehensible and convincing ways. Furthermore, and in contrast to other simulation studies, ABS studies *generate* the social dynamics “from the bottom up” (Epstein, 1999). That is, the social dynamics are not formulated as an aggregate in the model itself (for example as a differential equation) but, rather, arise from the modeled interactions of microelements (i.e., agents) in the simulation.

The wealth of detail in these models and their “bottom-up” approach are often taken as features that make agent-based models particularly suited for policy decision support. For example, it is often claimed that

decision makers might be more willing to trust findings based on rather detailed simulation models where they see a lot of economic structure they are familiar with than in general insights obtained in rather abstract mathematical models. (Dawid & Fagiolo, 2008, p. 354)

However, these very features are also the source of specific epistemic problems, creating an illusory familiarity. In particular, the greater detail and higher complexity of ABS increase the numbers of assumptions and parameters that need to be specified. This gives many more degrees of freedom to the ABS modeler and makes it difficult to restrict the ranges of model parameters based on empirical data. Consequently, the result of such simulation studies is likely not a unique model but an ensemble of models. This poses a problem for assessing the external validity of the models, which is already aggravated by their increased specificity. Furthermore, it requires new methods for making policy decisions based on such ensembles.

This article discusses the use of ABS for policy decision making at the hand of five simulation studies of smallpox epidemics. I focus on the smallpox case for the following reasons. First, it offers clear and well-understood decision alternatives. Second, the policy 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 ABS. Last, the studies are somewhat extreme in that empirical data on their subject is very scarce, thus allowing us to identify the issues of external validity particularly clearly.

The discussion of these case studies is supposed to serve three purposes.

1. I warn against taking “familiarity” with specific ABS as a criterion for having confidence in the model’s policy recommendations.
2. I show that specific epistemic issues—in particular the high number of detailed simulated systems—require additional reflection on which decision rules to choose for policy decisions based on ABS.
3. I point out directions in which I believe the construction and uses of ABS in policy decision could be improved.

The article is structured as follows: The second section introduces the specific policy issues regarding smallpox epidemics and distinguishes two modeling styles that may support policy decision making in this context. The third section surveys five agent-based smallpox models. The fourth section discusses how rational policy decisions can be made on the basis of these simulations. In particular, it is argued that standard expected-utility rules cannot be applied and that nonprobabilistic rules such as maximin or satisficing are sensitive to the scope of possible models. In the fifth section, I therefore suggest that the construction of ABS be theory guided to increase the systematicity of decision making based on them. The conclusion completes the article.

Modeling Epidemics

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. However, a significant minority of individuals experience severe (and often fatal) adverse effects to the vaccine, such that a national preventive mass vaccination (MV) may cause more deaths than an isolated smallpox epidemic (Kemper, Davis, & Freed, 2002). To make a rational policy decision, more information is needed about the potential development of an epidemic and the effects of different vaccination policies.

Unfortunately, the available empirical data are limited. Smallpox outbreaks have not occurred in the Western world since 1972; the virus has been eradicated globally since 1977. Most of the World Health Organization's policies, which led to its eradication, were employed in societies very different from those of the industrialized world. Differences in mobility, health infrastructure, and medical education made these cases difficult to compare with potential outbreaks in industrialized countries. As a remedy to this lack of relevant empirical data, scientists have proposed to study the effects of different vaccine policies on *computer-simulated* epidemics. Recent simulation attempts have focused on ABS studies to simulate smallpox epidemics and assess policy options (Brouwers, Mäkilä, & Camitz, 2006; Burke et al., 2006; Eidelson & Lustick, 2004; Epstein, Cummings, Chakravarthy, Singa, & Burke, 2004; Eubank et al., 2004; Halloran et al., 2002). 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.

The smallpox policy interventions most prominently discussed are MV and trace vaccination (TV). An 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 (postrelease MV). In TV, also called ring vaccination or targeted vaccination, every contact of a confirmed smallpox case is traced and vaccinated.¹

These policy options were first discussed on the basis of standard, non-agent-based epidemiological simulations. An often-cited example is Kaplan et al.'s (2002), which simulates an attack of 1,000 initial smallpox cases on a population of 10 million. The

population is assumed to mix homogeneously—that is, to consist of identical individuals, who have an equal chance of interacting with any other population member. R_0 , the rate of infections that 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 he or she vaccine sensitive; only in the third and fourth is he or she infectious; in the fourth, however, he or she shows symptoms (scabs) and is automatically isolated. Additionally, administration of vaccination is modeled under logistical constraints: MV of the whole population is achieved in 10 days. Tracking *and* vaccinating an infected person in TV, however, takes four times as many nurse-hours as simple vaccination.

Kaplan et al. (2002) thus offer an example of a classical equation-based simulation study. By assuming homogeneous mixing, the infection rate R_0 becomes a parameter characterizing the population. Policy effects then are modeled directly on this population parameter, and the main question is whether vaccine administration can outpace the random spread in the population. Unsurprisingly, maybe, the results heavily favor MV over TV. Both initiated on Day 5 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 the size of initial attack and changes in R_0 , further supporting the strong results in favor of MV. The time to identify and then vaccinate the exposed is simply too long for the specified R_0 and for the assumed period in which the exposed are still sensitive to the vaccine.

Models like this, which assume homogenous population mixing, have been criticized 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 in a population that matter but the patterns of exposure. These patterns are constituted by the social interactions of infectious and susceptible individuals. Explicitly modeling 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 & Lynch, 1999, p. 1170)

The so-called small-world effects 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 neighbors 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 characterized by $p = 0$, and a completely disordered network (similar to the case of homogeneous mixing) is characterized by $p = 1$. Kuperman and Abramson (2001) simulated the transmission of a (nonfatal) 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 matters greatly.

Results from Kuperman and Abramson (2001) and others make an important contribution to the vaccination policy discussion. They show that results such as Kaplan et al.'s (2002) must be taken with caution, because there may be causal factors at work that disturb the stability assumed in those models. Specifically, while the Kaplan et al. article claims that MV is advantageous over TV in *all* relevant epidemic contexts, the Kuperman and Abramson (2001) article cautions that in networks with a high degree of disorder, the occurring oscillations in the evolution of the epidemic may disturb the advantageousness of MV in some way. Thus the Kuperman and Abramson article can be used as evidence against the conclusiveness of models such as Kaplan et al.'s (2002; for further analysis of such models, see Grüne-Yanoff, 2009). Yet they do not themselves offer conclusions immediately useful for making policy decisions. Instead, by identifying causal tendencies that potentially contravene the results of those models, they make obvious the need for further investigation before a well-founded policy decision can be made.

Such further investigations can take various directions. Which direction is taken depends on the *modeling style* that is preferred by the modeler and the discipline in which he or she works. For example, one may investigate the dynamics of historical smallpox epidemics or the dynamics of other epidemics that are in relevant ways similar to smallpox. Such investigations may provide evidence whether the small-world effect is indeed relevant for vaccination policies. This modeling style is often found in economics, where a possible causal factor is isolated in a theoretical model (see Cartwright, 1989, Mäki, 2009) and the relevance of this possible factor for a particular context tested in an empirical model (by regression analysis on a particular data set). Should the investigation yield evidence that the small-world effect is not relevant in these contexts, it may give policy makers reason to make their decision based on the Kaplan model. Should it yield evidence for its relevance, however, it may also give further clues what the conditions are under which it becomes efficacious. These results can then also be included in the decision making process.

However, as I mentioned above, the empirical data on smallpox or similar epidemics is very limited.² Testing the efficaciousness of an isolated factor in a particular context, without access to relevant empirical data, is difficult. This may have been one reason why all the articles discussed here chose a different modeling style. Instead of investigating the relevance of the possible causes investigated in the theoretical models, they fill the theoretical model with more detail and investigate the result of these enriched models. Figuratively speaking, they take a model of the kind Kuperman and Abramson (2001) presented and *concretize* the abstract network into something that recognizably resembles a social structure. These attempts vary in detail, as I will show in the next section. The general purpose of this modeling style is to increase the amount of detail in tendency models in order to learn how the tendency plays out in the specific circumstances. Computer simulations here are of help, as they allow a massive inclusion of detail in the model.

This raises the question how such a concretization aids policy decision making. ABS offer the possibility to study the small-worlds effect in detail-rich settings, with a large degree of “familiar” structure replicated in the simulation. When a smallpox epidemic is simulated in these settings, and different vaccination policies tested, the simulations show how many people die, how high the costs are, how fast the epidemic disappears (in contrast to, e.g., Kaplan et al., 2002, who replace much structure and detail with stylized assumptions of homogeneous mixing). This detail and attention to “familiar” structure may help convince a policy maker of the significance of the model’s recommendations.

Yet the possibility of detail richness and “familiar” structure also has a reverse side. The more detail is included in the model, the more effort the modelers must make to justify the *external validity* of the model. This is an issue for all models, but because of their extra detail and structure, it is a particularly pressing issue for ABS. Principally, three questions arise:

1. In which ways does the ABS model relate to the real-world situation that the policy maker is interested in? Unlike with highly stylized models, the relation between ABS and world cannot be an abstraction relation. The ABS model presents a concrete simulated world; hence the relation is between two particulars.
2. Where differences between model and world appear, how does the policy maker justify decisions based on a model? Because of the massive detail in the ABS, it is implausible that it is an isolation relation, as sometimes argued for stylized models. Isolation simplifies the context and neutralizes disturbing factors (Mäki, 2005), while an ABS tries to replicate the complexity of the environment and include many different factors.
3. Where relations between model and world are unclear, how does the policy maker justify decisions based on a model? ABS is analytically not tractable and hence lacks the clarity needed for conceptual exploration, a common function for stylized models. Instead, as I will discuss in the fourth and fifth sections, ABS studies are sometimes used to map the space of possible futures.

Several smallpox ABS models have been developed. They build the epidemic macro phenomenon “from the bottom up”—by designing agents with heterogeneous attributes and heterogeneous behavioral rules interacting in discrete time in an explicitly modeled landscape. Five such models are surveyed in the next section.

Five Smallpox Models

Model 1

Halloran et al. (2002) simulate an attack of 10 infected agents on a stochastically generated community of 2000.³ The community consists of four neighborhoods, each containing large day care centers and small playgroups. Two neighborhoods share an elementary

school, and all four share a middle and a high school. Households in each neighborhood contain up to seven people. Agents are characterized by age, family ID, disease status, and vaccination status. Household size and age distributions are based on the U.S. 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 is thus 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. (2002) first derive a general household transmission rate from African/Asian data of the 1960s and 1970s. On that basis, they then derive relative magnitudes for adult-adult, adult-child, and similar 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. Two hundred simulations were performed for each intervention. At 80% MV after the first confirmed case, on average 0.9 out of 1,000 individuals died, while at 80% TV, 10.9 out of 1,000 died. Similar differences were found if vaccinations were administered after the 15th (9.4 vs. 19.6 deaths) and the 25th case (13.7 vs. 28.2 deaths). However, while MV prevented 0.5 cases per dose given, TV prevented 2.01 (both administered after the first case).

Model 2

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 neighbors (the so-called Moore neighborhood). 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 neighbors: 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 neighborhood. Homogeneous mixing is thus completely eschewed; instead, agents interact in a number of dynamic neighborhoods.

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 or she has, and whether the susceptible agent has partial immunity.

Burke et al. (2006) 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. Postrelease 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.

Model 3

Eidelson and Lustick (2004) simulate a single infected agent attack on a population of 1,764. 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 characterized by its coordinate address, its identity repertoire (a finite set of affiliations) and its currently activated identity. At each time step, identities change in response to neighbors' identities. If sufficient neighbors 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 his or her repertoire but above a certain threshold, he or she will adopt the identity into the repertoire. If it is above a higher threshold, he or she will switch to this identity at the same time.

Transmissions occur between direct neighbors. Transmission probabilities depend not only 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 removed from the population.

Simulation results show that TV rivaled MV if a very substantial proportion of small-pox 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, approximately 3.5% of the population is infected, while at 80% TV, approximately 6% is infected.

Model 4

Eubank et al. (2004) simulate an attack of 1,000 infected agents on the population of Portland, Oregon, of 1.5 million. Portland is represented by approximately 181,000 locations, each associated with a specific activity, such as work, shopping, school, and maximal occupancies. Each agent is characterized 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 U.S. census data.

Smallpox is modeled by a single parameter, disease “load” (analogous to a viral titer). Agents have individual thresholds, above which their load leads to infection (and load growth at individual growth rates), symptoms, infectiousness, and death. Every hour, infectious agents shed a fixed fraction of their load to the local environment. 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 24 hours 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.

Model 5

Brouwers et al. (2006) simulate an attack of 50 infected individuals on the population of Stockholm plus 1/10th of the Swedish population. The model is based on a 100×100 -meter grid of Sweden, each cell of which contains households, kindergartens, schools, offices, and hospitals according to the 1998 census. Agents are characterized 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. (2006) 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. 6,897,378) and that therefore TV is preferable to MV.

All the surveyed articles predict results for MV and TV that differ much less drastically than the results from Kaplan et al. (2002). How much the two policies differ in terms of average death rates—and hence how cost-intensive the prevention of additional deaths is—depends on the details of the respective models.

The studies obviously differ in the details of their model. More significantly, they also differ in the way these details are determined and (implicitly) justified. Studies 1 (Halloran et al., 2002) and 2 (Burke et al., 2006) rely on intuitive notions of “households,” “neighborhoods,” “towns,” and so on, to furnish the models. Study 3 (Eidelson & Lustick, 2004) appeals to a theory of agent identity formation that regulates social interaction in the model. This theory is empirically not well confirmed; as I will discuss in the fifth, it systematizes the model construction but does not justify it. Studies 4 (Eubank et al., 2004) and 5 (Brouwers et al., 2006) rely on large amounts of empirical data to determine many details of the model. Yet even these studies have to make many assumptions about model details that are not supported by empirical evidence and thus revert to appeals to intuition. The external validity of each model is thus a serious issue. How this issue affects the models’ use in policy decision making I discuss in the next section.

Making Decisions With ABS

The surveyed articles give more or less unconditional policy advice. To quote just two examples, “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, p. 1148); and “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, p. 180). These recommended measures evidently refer to interventions taken by real policy makers, not to manipulations of 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?

The surveyed ABSs demonstrate—in minute detail—the effects of diverse policy measures on otherwise identical worlds. It thus presents policy interventions in the classical form of a decision problem (Luce & Raiffa, 1957, p. 275): a list of possible alternative policy actions, each associated with either a single effect or a list of possible alternative effects. The policy maker then just needs to rank these effects, find a way to weight effects if more than one effect is associated with a policy, and derive the ranking of policy actions from this ranking of effects.

In the simplest and most ideal possible scenario, the ABS successfully imitates the relevant causal processes of the world in all its details and thus can be treated as a surrogate for the real system. One may be tempted to interpret Eubank et al. (2002) this way: The model they offer 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. If the model successfully imitated the relevant processes, the simulation results could be interpreted as showing how the real world would change under the simulated policy interventions. Furthermore, because the model they present is deterministic, it associates only one effect (i.e., the description of one final world) with each policy intervention. The policy maker would then face a decision under certainty: Each policy action would lead to one effect; each effect would be ranked (e.g., according to number of deaths); and the ranking of effects would trivially imply a ranking of policy interventions. Eubank’s et al. policy conclusions seem to be based on such a view: They recommend TV over MV based on the results that their ABS yields.

However, this interpretation turns out to be unwarranted. At closer inspection, it becomes clear that many model assumptions are unjustified. 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). Location occupancy rates crucially influence the number of possible contacts (and hence may be relevant for the disease spread). Furthermore, the contact data on which the simulation is based give a detailed account of social interaction only at one point in time. The model lacks any account how these social contacts may change under external shocks. The arrival of a threatening epidemic is, arguably, such a shock: It may well have important influence on how often people appear in public, go to work, or go to the hospital. Assuming that social interactions remain stable under such conditions is questionable.

Concerns about the correctness of modeling assumptions in ABS differ in two ways from concerns about idealization assumption in stylized, analytically tractable models. First, such models make idealizing assumptions for the purpose of simplicity, understanding, and computational tractability. That is, they make assumptions that may be false of

the target systems they are modeling, yet these assumptions facilitate the comprehension of what is modeled as well as the application of analytic solutions techniques. In contrast, ABSs are not in need of simplicity or tractability. The computer implementation implies that even very complicated models can be represented in comprehensible ways and that even very complex dependencies can be computed.⁴ Instead, ABSs face the problem that it is very difficult to determine whether all the detail necessary to run the simulations is justified or not. The need for detail thus raises the issue of uncertainty about the modeling assumptions, and this differs notably from the need for idealization in standard models.

Second, idealizing assumptions in standard models can often be justified. For example, it can be argued that the idealized parts of the model were never intended as representations of the world but instead were used to create a “neutral” environment in which the workings of some factor could be studied without impediment (see Mäki, 2005). Or, it could be argued that the idealizing assumptions do not affect the model result in relevant ways, in comparison with less simple and less tractable alternative assumptions. These justifications require that the processes through which the assumptions influence the model results are open to inspection and verification by the model user. Unfortunately, this is not the case with ABS, as these simulations are “epistemically opaque” (Humphreys, 2004, p. 147). Even a modeler with a very good understanding of a particular ABS will typically not be able to say how a specific assumption (say, of a behavioral rule) will affect the simulation outcome. ABS, after all, was designed to deal with high degrees of computational complexity. Thus, while the complex dependencies are programmable in an ABS, the effect of adjusting a certain assumption can typically not be foreseen but must be found out through renewed simulation runs. This increased dependency on the model makes it harder to justify the assumptions underlying ABS.

The need for detailed specifications, combined with uncertainty about the correct assumptions and the epistemic opacity of the derivation process, force decision makers to use ABS in a different way than they use standard models. It is unlikely that ABS modelers will be able to successfully argue for a *unique* specification of every micro-level detail. Empirical data will only act as *filters* of possible models. Observed macro patterns will exclude those simulations that fail to generate these patterns; and knowledge about micro properties will exclude those simulations that diverge from these properties. However, as I argued above, there are no data about actual smallpox epidemics in industrialized societies; hence pattern matching largely fails as a filter for the simulations discussed here.

What is left is the possibility to specify as many micro properties as possible and vary those properties whose specification cannot be justified across different simulations. Instead of using one basic model, and investigating the effects of varying policy interventions on it, the decision maker is forced to investigate these effects on a set or *ensemble* of simulation studies. Using such ensembles in decision making has been dubbed *exploratory modelling*.⁵

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, p. 435)

Policy makers may use such ensembles of models in different ways. In the simplest case, a probability distribution is specified over the set of models. If each model specification had a determined probability assigned, then the standard decision procedure under risk, expected utility maximization, could be applied. The policy maker could identify the outcomes of each vaccination policy, as obtained from different model specifications; he or 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). Finally, he or 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 ranking of policy, according to which the policy maker could justify his or her decisions.

However, for any of these models, specifying such a probability distribution would be difficult. The models' probability cannot be identified directly. The simulated worlds are so detail rich and complicated that one cannot reasonably argue for a probability assignment *simpliciter*. Instead, such assignments must be derived from the models' elements.⁶ Therefore, disregarding dependence effects, the models' elements must have probabilities assigned, such that the probability of the whole can be derived from them. For many of the discussed models, some of the specific attribute values for each individual and each interaction space can be drawn from an empirically supported distribution (e.g., census data). However, not all elements of the model thus receive a probability assignment. In particular, no such data are available for the agents' behavioral rules. This is, first, because of a lack of data on "how people interact" on such a level of generality. Second, this is because the probability distributions over such behavioral rules are likely to be conditional on all sorts of things, such as environments, the frequency of others' behavioral rules, and external shock. It would require an empirically well-supported theory to justify the assignment of such conditional probabilities, and no such theory is currently available.

Without justified probability assignments, the method of explanatory modeling must explain how a decision can be reached on a nonweighted set of models. The lack of probabilistically weighted models makes simulations yield only a "landscape of plausible futures" (Banks, Lempert, & Popper, 2001, p. 73). How can the policy maker base his or 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" (Banks, 1993, p. 440). Each plausible model assigns an outcome to each vaccination policy. Instead of comparing the policies in all outcomes, the policy maker may "hedge" and only compare the policies with respect to their worst outcomes. The policy maker thus justifies his or her decision with reference to the *maximin* procedure: he or she chooses the policy that maximizes the minimal (worst) outcome.

An alternative approach would pay equal attention to all models and choose the policy that performs relatively well, compared with the alternatives, across the range of plausible futures. This robustness analysis of policies (Lempert, 2002) can be interpreted in two ways.⁷ If "performs relatively well" means that a certain policy *P* performs better than

other policies for *each* plausible future—that is, 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 nondominated 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*. Such a decision procedure specifies minimal thresholds for simulated policy interventions. For example, a policy may be deemed acceptable if it yields (simulated) results above a certain threshold in at least two thirds of the cases. Thresholds must be argued for in terms of the specifics of the policy goals, and they can be refined in such a way that they are sufficiently selective.

Both maximin and satisficing are very sensitive to the number of models considered. The wider the scope, the more likely the inclusion of some outlandish terrible future, which will affect maximin choice. Similarly, the wider the scope, the more likely the inclusion of some outlier below the threshold, which will affect satisficing choice. Given the uncertain status of many model specifications, exploratory modeling is prone to such misspecifications.

Selecting the Model Ensemble

Because the scope of the model ensemble so crucially influences the decisions based on it, the question arises how this scope is determined. As I argued in the previous sections, all discussed models are seriously underdetermined by the empirical evidence. Consequently, all surveyed studies must specify values of empirically undetermined parameters and behavioral rules. Yet by what strategies these values are set influences the scope of the model ensemble and hence the policy decision. In particular, these strategies (a) determine which models go into the model ensemble and (bi) determine which models are excluded from the model ensemble.

Most of the studies discussed here rely on intuition to specify the empirically undetermined parameters of their models. Eubank et al. (2004) offer “reasonable guesses,” and Burke et al. (2006) appeal to intuitive beliefs:

We include the level of social detail that we believe necessary to capture the transmission dynamics of smallpox. (p. 1148)

Yet such intuitions pose general problems for scientific practice, and particular problems for ABS. The general problem is that intuitions are difficult to account for and are difficult to share intersubjectively. That Eubank et al. (2004) think their guesses “reasonable” and Burke et al. (2006) believe a certain level of detail to be relevant give only very limited support to these parameter determinations. One may claim that these are expert judgments; yet the reader is not presented with any evidence on which such judgments could be based and hence should remain skeptical. Furthermore, even if evidence was presented, the judgment itself is never accounted for. Because of the complicated nature of many modeling endeavors, such “intuitions” too easily become hidden receptacles for significant but dubious assumptions.

This leads to the problems specific to ABS. First, intuitions tend to unduly restrict the scope of the model ensemble. Cognitive biases such as representative heuristics or confirmation biases show that people tend to confirm their preconceptions. Even experts are subject to such biases (see Baron, 2000). Scientists constructing simulations with particular results in mind (in the cases discussed here, the superiority of TV over MV) are likely to find those intuitions plausible that confirm their expected results. Such biases are of course operative in all model building. Yet the complicated and detail rich nature of ABS makes it easier for such biases to remain undetected. The fact that an ensemble of models is constructed makes this detection even more difficult.

Second, one should be less confident about modelers' intuition concerning ABS than about simpler models. I have argued elsewhere (Grüne-Yanoff, 2009) that intuition is an important guide for constructing models of the traditional type. However, I also pointed out that intuition is formed at the specifics of the model world. One does not have independent intuitions about how people behave *in abstracto* but instead forms intuitions when confronted with the specifics of the world the model describes. This makes intuition a difficult basis for ABS. The large number of and the enormous details in the possible worlds simulated do not allow the kind of mental immersion that traditional modeling thrives on. ABSs, after all, are constructed to be processed by a computer and not by human minds. Judging ensembles of simulated worlds may simply go beyond human intuitive capacities.

Thus, intuitions are difficult to account for, prevent intersubjective standards, and are subject to cognitive biases. For ABS in particular, they are likely to restrict the model ensemble unduly or misspecify it altogether. For these reasons, intuition seems a nonoptimal strategy to determine and manage model ensembles.

Instead, I suggest the alternative strategy of *theory guidance*. Theories are sets of general principles that are intended to account for certain phenomena. Because of their generality, theories also constrain what is believed to be possible. They exclude certain simulated worlds, because these worlds conflict with their principles; and they mandate the inclusion of other worlds into the set of possible models, as these models conform to the theoretical principles.

In comparison with a strategy guided by intuitions, the theory guided strategy has five advantages.

1. It makes explicit the assumptions that are often hidden behind the appeal to intuitions. These assumptions thus become accountable.
2. A theory allows judging such assumptions in an intersubjectively agreeable way. Even if empirical confirmation is lacking, a theory can show how an assumption is derived from other assumptions and what the most basic assumptions look like.
3. A cognitive bias is more easily detectable in theories that make the basis of certain assumptions explicit than in intuitive judgments. To avoid criticism of such bias, it is likely that theory-guided ensemble construction will increase the ensemble's scope.
4. The explicit formulation of guiding theoretical principles allows for a more effective application of computational resources than an intuition-based strategy.

Intuitions mix structural and quantitative judgments, because human intuitions are biased toward concrete situations. In contrast, theoretical hypotheses can abstract from many concrete aspects of a situation. For example, they can determine only the abstract structure of all models in a model ensemble, leaving the parameter values open. This would allow for modelers to “sweep the parameter space” (Epstein, 1999, p. 51)—that is, compute the result for all possible parameter sets.

5. Theory-guided model ensembles are much easier to adjust when new information becomes available. The theoretical hypothesis that selected all members of the ensemble can be tested against the new information, and adjusted if needed; the adjustment of the ensemble in accordance with the hypothesis adjustment is then just a technicality. In contrast, adjusting an intuition-guided model ensemble in the light of new information is a much messier affair.

Of the cases studied, only Eidelson and Lustick (2004) derive the crucial behavioral assumptions in their model from a general theory. This theory—the agent-based identity repertoire model (Lustick, 2000)—is not well confirmed and therefore does not warrant the use of these assumptions. Nevertheless, exploratory modeling guided by agent-based identity repertoire is more accountable, intersubjectively accessible, less cognitively biased based, more compatible with computational methods, and more adjustable in the light of new information than modeling guided by intuitive judgment.

To promote theory guidance is an attempt to explicitly formulate a strategy for ABS modeling for policy decision purposes. Modelers should explicitly state (or develop) theoretical hypotheses first and then construct ABS in compliance with these hypotheses. Any maximin or satisficing decision taken on the basis of such a model ensemble will be much better understood than if the model ensemble is based on empirical data alone or based on intuitive judgments.

Conclusion

In discussing five cases of smallpox ABS studies, I argued that policy decision making faces particular and novel problems when based on ABS. In particular, the epistemic opaqueness of the simulation processes and the uncertainty about the simulations’ assumptions are issues that are more pressing for ABS than for traditional modeling styles. Consequently, in many cases—and in particular in the smallpox simulations discussed—policy decisions cannot reasonably be based on expected utility calculations from a single simulation. Instead, alternative, nonprobabilistic decision rules have to be applied to ensembles of models. Yet as I show, the exploratory modeling approach that propagates such model ensembles lacks a systematic strategy how to construct these ensembles. As a first step in the direction of a more systematic strategy, I propose that model ensembles be constructed in compliance with a theoretical hypothesis. While the chosen theoretical principles are unlikely to justify the model construction, they make the scope and limitations of the model ensemble explicit and hence systematize the construction process.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the authorship and/or publication of this article.

Funding

The author(s) thanks the Collegium of Advanced Study for supporting research on this article.

Notes

1. The Centers for Disease Control and Prevention (2002) 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.”
2. Small-world effects are a reason why data from other contagious diseases cannot easily be adopted for modeling smallpox epidemics. Common flue and others differ from smallpox in their incubation period, their modes of transmission, and the relation of infectiousness and symptomatic periods—that is, exactly in those properties that influence the causal relevance of the patterns of exposure.
3. An extension of the model for larger populations and with a more detailed disease natural history is presented in Longini et al. (2007).
4. Of course, ABSs are also constrained, namely, by the computing power available to current modelers. But the analytic tractability is a much stricter limitation than computing power (and one that is unlikely to go away).
5. Exploratory modeling is different from sensitivity analysis. Commonly, sensitivity analysis varies only one parameter at a time and records the effect of this variation on the outcome of the model. Exploratory modeling, in contrast, allows varying all parameters, also simultaneously, that are not constrained by our present knowledge. The set of models produced by sensitivity analysis is thus a proper subset of the set produced by exploratory modeling.
6. A Bayesian may argue that the derivation of a probability distribution over models from probability distributions over models’ elements is unnecessary. Instead, the modeler may just assign subjective probabilities and update these probabilities as he or she obtains new information. With repeated updating, so the Bayesian claim, different subjective probability assignments of different modelers will eventually converge. However, even if a Bayesian may assign subjective priors to the rules, he or she would lack a procedure by which one could assure that such subjective probability judgments will eventually converge. Specifically, we cannot simulate past data models of epidemics with the given ABM, because such epidemic data are lacking. Thus, the practice of updating subjective priors through this method is impossible for the kinds of simulations at hand. A justified assignment of probabilities to stochastic models of this sort therefore seems out of reach.
7. Note the difference between robustness analysis of policies (checking the value variation of different possible outcomes of a policy) and sensitivity analysis of a model (which is also often called robustness analysis).

References

- Bankes, S. (1993). Exploratory modelling for policy analysis. *Operations Research*, 41, 435-449.
- Bankes, S. C., Lempert, R. J., & Popper, S. W. (2001). Computer-assisted reasoning. *Computing in Science and Engineering*, 3, 71-77.
- Baron, J. (2000). *Thinking and deciding* (3rd ed.). New York, NY: Cambridge University Press.
- Brouwers, L., Mäkilä, K., & Camitz, M. (2006). *Spridning av smittkoppor—Simulerings experiment* [Spread of smallpox—Simulation experiment]. *SMI-Rapport*, 5, 2006T.
- Burke, D. S., Epstein, J. M., Cummings, D. A., Parker, J. I., Cline, K. C., Singa, R. M., & Chakravarty S. (2006). Individual-based computational modeling of smallpox epidemic control strategies. *Academic Emergency Medicine*, 13, 1142-1149.
- Cartwright, N. (1989). *Nature's capacities and their measurement*. New York, NY: Oxford University Press.
- Centers for Disease Control and Prevention. (2002). *Smallpox response plan and guidelines (Version 3.0), Guide B, Part 2*. Retrieved from <http://www.bt.cdc.gov/agent/smallpox/response-plan/files/guide-b-part2of3.pdf>
- Dawid, H., & Fagiolo, G. (2008). Editorial. *Journal of Economic Behaviour & Organization*, 67, 351-354.
- Eidelson, B. M., & Lustick, I. (2004). VIR-POX: An agent-based analysis of smallpox. Preparedness and response policy. *Journal of Artificial Societies and Social Simulation*, 7(3). Retrieved from <http://jasss.soc.surrey.ac.uk/7/3/6.html>
- Epstein, J.M. (1999). Agent-Based Computational Models and Generative Social Science. *Complexity*, 4(5), 41-60.
- Epstein, J. M., Cummings, D. A. T., Chakravarthy, S., Singa, R. M., & Burke, D. S. (2004). *Toward a containment strategy for smallpox bioterror: An individual based computational approach* [Brookings Monographs]. Washington, DC: Brookings Institution Press.
- 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.
- Grüne-Yanoff, T. (2009). Learning from minimal economic models. *Erkenntnis*, 70(1), 81-99
- Halloran, M.E., Longini, I.M, Nizam, A., and Yang, Y. (2002). Containing bioterrorist smallpox. *Science*, 298, 1428-32.
- Humphreys, P. (2004). *Extending Ourselves: Computational Science, Empiricism and Scientific Method*. Oxford: Oxford University Press.
- Kaplan, E. H., Craft, D. L., & Wein, L. M. (2002). Emergency response to a smallpox attack: The case for mass vaccination. *Proceedings of the National Academy of Sciences*, 99, 10935-10940.
- Kemper, A. R., Davis, M. M., & Freed, G. L. (2002) Expected adverse events in a mass smallpox vaccination campaign. *Effective Clinical Practice*, 5, 84-90.
- Koopman, J., & Lynch, J. (1999). Individual causal models and population system models in epidemiology. *American Journal of Public Health*, 89, 1170-1174.
- Kuperman, M., & Abramson, G. (2001). Small world effect in an epidemiological model. *Physical Review Letters*, 86, 2909-2912.

- Lempert, R. J. (2002). New decision sciences for complex systems. *Proceedings of the National Academy of Sciences*, 99(Suppl. 3), 7309-7313.
- Longini, I. M., Halloran, M. E., Nizam, A., Yang, Y., Xu, S., Burke, D. S., . . . Epstein, J. M. (2007). Containing a large bioterrorist smallpox attack: A computer simulation approach. *International Journal of Infectious Diseases*, 11, 98-108.
- Luce, R. D., & Raiffa, H. (1957). *Games and decisions*. New York, NY: Dover.
- Lustick, I. S. (2000). Agent-based modelling of collective identity: Testing constructivist theory. *Journal of Artificial Societies and Social Simulation*, 3(1). Retrieved from <http://jasss.soc.surrey.ac.uk/3/1/1.html>
- Mäki, U. (2005). Models are experiments, experiments are models. *Journal of Economic Methodology*, 12, 303-315.
- Mäki, U. (2009). MISSing the world. Models as isolations and credible surrogate systems. *Erkenntnis*, 70(1), 29-43.

Bio

Till Grüne-Yanoff is a fellow of the Collegium of Advanced Study at the University of Helsinki. Before that he held appointments at the Royal Institute of Technology, Stockholm, and the London School of Economics. His research focuses on the methodology of economic modeling, on decision and game theory, and on the notion of preference in the social sciences. He has published in journals such as *Synthese*, *Erkenntnis*, *Theoria*, and *Journal of Economic Methodology*; among others, and has edited (together with Sven Ove Hansson) a book, *Modelling Preference Change* (2009). Contact: till.grune@helsinki.fi.