Chapter 3
Applications of Agent-Based Modeling and Simulation to Healthcare Operations Management

Sean Barnes, Bruce Golden, and Stuart Price

1 Introduction

Simulation techniques have been applied to problems in healthcare operations management for many years in several focus areas. Historically, research in this area has consisted of results derived from system dynamics (SD) and discrete event simulation (DES) methods (Fone et al. 2003; Jun et al. 1999; Koelling and Schwandt 2005). Both methods focus on system-level behavior, but they differ in how the system is modeled and how time is simulated. SD models represent entities as continuous variables whose states change continuously with time, whereas DES models contain individual components whose states only change at discrete moments in time. In either case, the goal is to aggregate the system behavior and draw conclusions about how the system evolves over time under internal and external forces. These techniques can provide valuable insight to problems in healthcare operations management and are ideally suited for many such problems.

In recent years, a new modeling and simulation methodology has gained momentum with respect to healthcare applications. The methodology is most commonly known as agent-based, or individual-based, modeling (ABM). In contrast to SD and DES methodologies, ABM focuses on modeling individuals, interactions between individuals, and in some cases, interactions with a physical or influential surrounding environment (Macal and North 2007b). This activity can then be aggregated to simulate how a system behaves over time. The focus on detailed, individual agents and their interactions makes ABM an ideal tool for analyzing complex systems such as healthcare facilities and organizations because there

S. Barnes (✉) • B. Golden • S. Price
Robert H. Smith School of Business, University of Maryland, College Park, MD, USA
e-mail: sbarnes@rhsmith.umd.edu; bgolden@rhsmith.umd.edu; sprice@rhsmith.umd.edu

are many components to these systems and outcomes can be difficult to predict without adequate model representation. A comparison between the three modeling frameworks is shown in Table 3.1.

Agents can interact with each other in many ways. Interactions may occur in a spatial environment, which could be a simple one-dimensional ring, a two- or three-dimensional Cartesian grid, or a specific geographic location or region. Alternatively, interactions can lack spatial representation and instead be constrained by relational considerations, where agents only interact if there is an explicit connection between them, such as being a member of the same family, working together, or being cared for by the same healthcare worker. For these cases, physical space has no effect on the outcome of the interactions, and, therefore, it does not need to be modeled explicitly. Social network analysis (SNA) is often paired with ABM for problems in which the structure of the interaction network is not uniform, implying that each individual may only interact with a specific subset of the population. For these types of problems, agents are often represented by nodes in a network, and interactions are represented by edges, which can be weighted to represent the frequency or type of those interactions.

There are several advantages of ABM over SD and DES. First, ABM is a more realistic modeling approach for many problems, especially problems in which there are multiple types of actors that interact in different ways. For these cases, it is very straightforward to model these actors as agents that have distinct sets of behaviors and characteristics without making assumptions as to how the system would be affected by each type. Individuals are not typically represented in SD models, and in DES models, individuals are explicitly modeled, but their states are typically a function of their status in some type of predefined system process. In addition, agent-based models facilitate detailed analysis of both individual- and system-level behavior because metrics at each level can be updated with each interaction. Agent-based models are also easier to explain than most SD and DES models because of their direct correlation to reality, which is an important factor in gaining the confidence of healthcare professionals and ultimately having an impact. SD models often consist of mathematical models that can become quite complicated, and DES models are usually described by intricate flow diagrams. These abstractions can cause difficulty in explaining model concepts to healthcare professionals who are not trained in mathematical or computational disciplines.

As in all methodologies, there are disadvantages to ABM as well. ABMs can become very complex when they incorporate a lot of detail. When this happens, it becomes difficult to separate the actual effect of each input parameter in the model. In addition, agent-based models can become computationally expensive, requiring excessively long computer run times for simulations. This problem has been alleviated to some degree by high-performance and parallel computing techniques, but it demands additional developmental resources that are not typically required by SD or DES models. Agent-based models also face different challenges related to assumptions. SD and DES models often require more assumptions about system-level parameters (e.g., patient admission rates, average hand-hygiene compliance of healthcare workers), whereas ABM requires more assumptions about

Table 3.1 Comparison of system dynamics (SD), discrete event simulation (DES), and agent-based modeling (ABM) methodologies

Attribute	SD	DES	ABM
Model perspective	System level	System level	Individual level
Level of realism	Low	Moderate	High
Model flexibility	Low	Moderate	Very high
Time domain	Continuous	Discrete	Discrete
Run times	Very fast	Moderate	Slows with increasing system size and complexity
Model inputs	Rate parameters and flow characteristics	Entity types, arrival times, queuing parameters, resource scheduling, process flows	Agent characteristics and interactions, environment specification
Model outputs	Dynamics, steady-state values, analytic expressions	Wait times, resource utilization, throughput	Unlimited
Software	Spreadsheet (e.g., Microsoft Excel), any programming language or mathematical software (e.g., MATLAB[a], Mathematica[b])	Object-oriented programming languages (e.g., C++, Python[c], Java), simulation software packages (e.g., Arena[d], AnyLogi[e], SIMUL8[f]), or mathematical software (e.g., MATLAB, Mathematica)	Object-oriented programming language (e.g., C++, Python, Java), open-source ABM software (e.g., NetLogo[g], Repast[h], MASON[i], Swarm)

[a] http://www.mathworks.com/products/matlab/
[b] http://www.wolfram.com/mathematica/
[c] http://www.python.org
[d] http://www.arenasimulation.com/
[e] http://www.xjtek.com/
[f] http://www.simul8.com/
[g] http://ccl.northwestern.edu/netlogo/
[h] http://repast.sourceforge.net/
[i] http://cs.gmu.edu/~eclab/projects/mason/

Table 3.2 Agent characteristics and definitions

Characteristic	Definition
Autonomy	Agents act independently of other agents
Heterogeneity	Agent characteristics and evolution of state are sufficiently different for all agents
Awareness	Agents can have varying levels of knowledge of the system state, ranging from ignorance to omniscience
Memory	An agent can remember its state and/or the state of the system at earlier points in time
Adaptation	Agents can change their behavior over time based on the current state of the system or prior experience
Goal oriented	Agents act toward accomplishing an objective
Rationality	Agents act in their best interests
Interactivity	Agents can exchange information or resources with other agents
Reactivity	Agent state or behavior can change in reaction to the environment or changes in the behavior of other agents
Mobility	Agents can move within the environment

individual-level parameters (e.g., probability of nurse-to-patient transmission for a given infectious disease) and the nature of interactions. Some of these requirements can be satisfied easily, by speaking with experts who have experience working in these environments, but others can be more difficult to quantify. However, the advantage of ABM is that few assumptions need to be made about the system as a whole, because system behavior is determined by the activities at the individual level.

Agents have several characteristics that can be used to distinguish ABMs from SD or DES models. These characteristics and their definitions are summarized in Table 3.2. For a more complete discussion on ABM, the reader is referred to Macal and North (2007a). There are additional characteristics associated with agent-based models, but this set is sufficient for identifying appropriate studies. Agent-based models do not necessarily possess all of these characteristics, and in some cases, the distinction between ABM and DES can be difficult. For the purposes of this review, each included research article must present the results of a simulation that models sufficiently heterogeneous agents that interact in a dynamic, nondeterministic manner. In general, articles that only describe the concept and implementation of an agent-based model without demonstrating insightful results were also excluded.

As with any modeling methodology, choosing an appropriate software package is critical during the early stages of developing an agent-based model. This selection process for healthcare applications depends primarily on three factors: the complexity of the model, the skill of the developer, and the level of interaction the consumer wants to have with the model. Developers with limited modeling experience and enthusiasm for modeling should identify the platform that is easiest to learn and provides the necessary functionality. More experienced developers are probably either skilled in a programming language that offers the required capability

or willing to learn a new language if it is more appropriate to the nature of the model. On the delivery end, some consumers may only require the results from a model and do not need to interact with it in any way. In this case, visualization and animation would be helpful capabilities, but a graphic user interface would not be required. If the consumer will interact with the model regularly, a graphic user interface is necessary and proper training needs to be provided.

Software can range from commercial products to a variety of programming languages and open-source platforms. Commercial software is typically easy to learn and is accompanied by extensive documentation and customer support, but it may not provide the modeling flexibility necessary for modeling a complex system. Object-oriented programming languages are especially well suited for agent-based models because objects correspond well to agents and object methods correspond to agent behavior, which facilitates a relatively straightforward implementation. Mathematics-based software typically provides substantial built-in functionality, but computational times can be prohibitively slow for an agent-based model. Lower level programming languages typically execute significantly faster than the aforementioned software types, which may be critical for large models, but they lack the built-in functionality that helps to accelerate development times. There is a growing set of dedicated software packages that offer built-in functionality common to agent-based models. These tools vary significantly in terms of their user-friendliness, level of documentation, speed, and modeling flexibility, but their diversity demonstrates the variety of applications that can be modeled using agent-based methods.

In this chapter, we review and evaluate a selected body of research that has applied ABM and simulation techniques to healthcare operations management, which we interpret broadly. Specifically, research that applies agent-based methods in the following areas is included, based on surveys of simulation and system dynamics applications to healthcare (Fone et al. 2003; Jun et al. 1999; Koelling and Schwandt 2005):

- *Healthcare delivery*: studies that focus on a single facility (e.g., a hospital) or unit within a department (e.g., an emergency department) and emphasize patient flow characteristics, scheduling, or allocation of resources.
- *Epidemiology*: studies that focus on the spread of illness or disease or the physiological understanding of an illness or disease.
- *Healthcare economics and policy*: studies that focus on the financial and policy decisions made by hospitals or healthcare organizations which can include the purchase of equipment, pharmaceutical items, or other medical equipment.

We review the ABM literature for each topic and briefly describe the core methods, summarize the key results, and identify best practices. We will also highlight areas within each topic where ABM and simulation filled a significant gap that was not addressed previously by other methods. Finally, we will propose some new questions that may be of interest moving forward.

2   Healthcare Delivery

There are many different types of facilities that serve as healthcare providers. Each provider has its own set of objectives and constraints. Outpatient facilities tend to operate using appointments that are scheduled well in advance, whereas emergency rooms deal almost entirely with unscheduled arrivals. This section focuses on modeling hospitals and the individual departments therein. There are models of other facilities such as long-term care facilities, outpatient clinics, or diagnostics laboratories, but the literature is not as extensive as it is on hospitals.

Hospitals are complex service systems that operate under many constraints. Providing healthcare services to patients is their primary objective, but there are many obstacles to doing so effectively. First, almost every patient has a different set of medical needs. Therefore, hospitals need to have a robust strategy to accommodate their needs. In addition, some patients are scheduled, whereas other patients (e.g., emergency and ambulance patients) enter the hospital unexpectedly. Failure to meet these demands could adversely affect the health of the patients, a far greater consequence than just lost revenue.

Hospitals have to manage many resources—including physical and human—in order to meet the needs of the patients they serve. They employ physicians, nurses, technicians, administrative personnel, and many other support staff that help to provide the necessary care to patients. Employees must be scheduled in such a way as to properly staff the hospital in the presence of a highly variable workload. On the other hand, these personnel are limited in the number of hours they can work before patient safety becomes a concern. Medical equipment can also be a critical resource. Some equipment is expensive and must be employed at high utilization rates in order to be cost effective. However, shortages in availability can cause bottlenecks in the system, which lead to excess waiting times for patients. Ultimately, the goal is to maximize the utilization of medical personnel and equipment while maintaining a sufficient buffer to accommodate spikes in patient demands.

A substantial amount of work has been done using DES to model the delivery of healthcare in hospitals. DES is a natural extension to the traditional queuing theory models, which are popular analytic models for service-oriented problems. DES allows for the use of distributions of system parameters that might be intractable in a purely analytic model but better fit observed data. For a survey of DES in hospitals and healthcare clinics, see Jun et al. (1999). DES is the current standard, and some of the agent-based models presented have yet to take full advantage of their potential in providing unique insight into these types of problems. In addition, these preliminary models have the potential to address questions inaccessible to DES. ABM can generate heterogeneous patients that each have a unique set of medical needs. Patient arrival rates can be a function of time or other system parameters. Moreover, dynamic responses by the physicians and nurses in the system can be tested, allowing for them to respond to not only the needs of each patient but also the state of the entire system. In addition, ABM can represent many resource types, including medical equipment, to analyze the system efficiency under various

patient demands and constraints. The following two sections deal with patient flow through the operating room and the emergency department, respectively. The third section examines the use of radio-frequency identification (RFID) technology in an emergency department.

2.1   Patient Flow

With rising costs and an aging population, the demand for affordable healthcare is likely to increase significantly. In the USA, the Bureau of Labor Statistics predicts that the healthcare industry will generate more new jobs than any other industry between 2008 and 2018, but increasing the workforce is not enough and does little to help control expenses. Beyond increasing the capacity of the healthcare system, it is important that facilities are operated efficiently to minimize the idle time of operating rooms, imaging equipment, and healthcare workers.

Workflow problems are often modeled using queuing theory (see Chap. 1). Queuing theory facilitates analysis of complex systems that serve individuals with either stochastic arrival times, service times, or both. For a stationary arrival process, it is straightforward to determine the number of servers that can adequately meet the system demands. However, when the arrival process is less predictable, as it is in hospitals where patient arrival times typically vary by time of day and day of week, analytic solutions for these problems are not tractable. As a result, simulation becomes the best approach to evaluate potential solutions for improving the system throughput.

One example is a simulation model developed by Pearce et al. (2010), in which the authors seek to minimize delays in the first set of morning surgeries in an operating room (OR). They focus on delays at the start of a period because these delays can cause successively longer delays throughout the day. Patients scheduled for surgery are either inpatients or outpatients. All outpatients must first pass through registration and then receive some subset of tasks (e.g., triage, phlebotomy, etc.) before they join the inpatient track. Inpatients and processed outpatients see a combination of healthcare workers, which may include a patient care technician, surgeon, anesthesiologist, and certified registered nurse anesthetist before finishing their preoperative process. The tasks in both the initial outpatient track and the subsequent inpatient track may be subject to reordering in order to minimize wait times and maximize resource utilization. After finishing both tracks, the patient is then queued for the operating room.

The model developed by Pearce et al. treats patients, staff, and informational units (e.g., charts and laboratory results) as agents. Patients progress through four states: waiting room, preoperative room 1 (for the outpatient track), preoperative room 2 (for the inpatient track), and the OR itself. Hospital records are used to determine relative inpatient/outpatient loads. Data for the time each task requires was gathered by several days of shadowing patients through the process.

The expected length of preparatory time is a function of acuity level, with sicker patients taking more preparatory time on average than patients with lower acuity levels.

A variety of tests are run to improve the percentage of initial surgeries that began within 10 minutes of the scheduled time while keeping patient wait times as short as possible. Several scheduling strategies are evaluated. One strategy that was found to be effective is to schedule high-acuity patients and patients requiring blood work as early as possible because they have longer service times on average and more variance. The best improvement is achieved by signaling and coordinating the agents that needed to see the patient upon arrival to preoperative room 2. This strategy results in an increase of on-time surgeries from a baseline of about 60% to one of 85%. Signaling also allows the benefits of other changes to be more fully realized. For example, a policy of using the RN staff to perform triage without signaling results in only about a 1% improvement in on-time surgeries. Without signaling, patients under this policy spend less time in the waiting room but more time waiting in the preoperative stage. When signaling is used in addition to this policy, it allows for around an 89% on-time rate, about a 4% improvement over the base case with signaling. Future work will focus on resource utilization at each of the tasks for the inpatient and outpatient track.
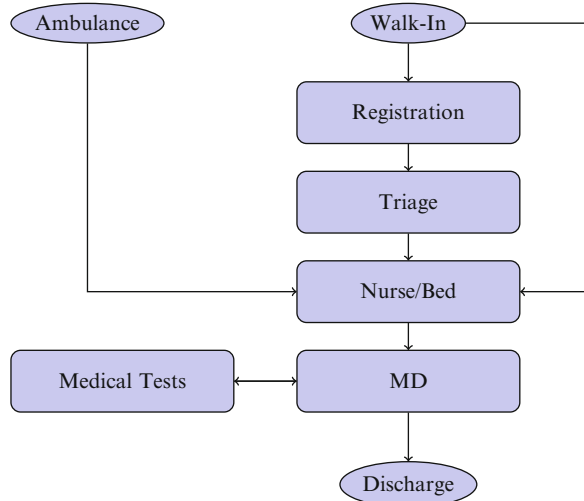
2.2 Emergency Departments

Emergency rooms have become a primary source of medical treatment for an increasing number of people. According to the CDC there were over 123 million visits to emergency departments in 2008, a 37% increase from 1996. Emergency rooms, more so than other parts of the hospital, have a highly variable workload because patients by their very nature are unscheduled. Therefore, one of the key problems is how to appropriately staff physicians and nurses in order to have a robust response to uncertain demand. Such a response, ideally, has neither excessive patient waiting nor idle physicians or nurses. Another key question is identifying bottlenecks in the system. For example, what should be the relative ratio of staffing of nurses to physicians? In addition to staffing, what procedures can be implemented to reduce waiting times? This might include steps taken within a single department, such as criteria for adding or subtracting triage nurses based on queue length. On a regional scale with multiple hospitals, how can ambulances and patients be diverted to prevent excessive congestion at any one hospital in the system?

Models of patient flow through an emergency department (ED) typically simulate a process similar to that in Fig. 3.1. The patients enter the system as either walk-ins or arrive by ambulance. Patients who arrive via ambulance circumvent the registration process and proceed directly to treatment as they have, in a sense, already been assigned a critical diagnosis. Walk-in patients first queue to register and then, after registration, proceed to queue for triage. At triage, a nurse will assign the patient a priority based on their emergency severity index (ESI). While the number of levels in an ESI need only be two or greater, the trend is toward a five-tiered

Fig. 3.1 The basic workflow of patients through an emergency department. Note that queues may form prior to any rectangular station in the system

Table 3.3 Summary of emergency department model parameters

Model parameters	Kanagarajah et al. (2006)	Jones and Evans (2008)	Laskowski et al. (2009)	Wang (2009)
Software platform	Micro Saint Sharp	NetLogo	C++	NetLogo
Adaptive agents	Yes	No	No	Yes
Medical tests	Yes	No	Yes	Yes
Ambulance arrivals	No	Yes	Yes	No
Triage levels	3	5	3	1

system. ESI in a five-tiered system ranges from 1, typically labeled resuscitation which must be treated immediately, to 5 for nonurgent patients. Patients are then queued separately with those receiving the highest ESI typically being served first (as opposed to a purely FIFO system). A certain percentage of walk-in patients are treated as ambulance patients (in that they are visibly in such need that they skip the registration and triage processes and proceed directly to a room or bed to await physician care). After triage, patients queue for a room or a bed where they will then queue for treatment by a physician. Some models have a physician order tests for a percentage of patients, the results of which will be reviewed by a physician before discharging the patient. Patients can exit the ED at any step in the process if they leave against medical advice, leave before receiving treatment, or die. Patients complete the system after seeing a physician when they are either discharged from the hospital, transferred to another hospital, or admitted to the hospital. A summary of models using this basic framework can be seen in Table 3.3.

The Kanagarajah et al. model focused on balancing economic incentives, workload, and quality of care (Kanagarajah et al. 2006). These competing goals ultimately are a function of physician utilization and wait time. The simulation

tests a variety of effects which included using on-call physicians to aid when arrival rates surged, adjusting the capacity of examination rooms where physicians serviced triaged patients, and having physicians spend less time with patients when lines were long. The amount of time spent with patients by a physician is an example of an adaptive behavior, which is unique to each individual physician and evolves over time. This adaptive behavior is a feature of ABM that cannot be easily replicated by other modeling techniques.

The Jones and Evans model, which allowed for a high degree of control in terms of inputs, was validated using data gathered from an urban hospital's ED (Jones and Evans 2008). Using a heterogeneous Poisson arrival process for patients, the authors observed how hourly and daily variation in patient arrival rates coupled with staffing levels affects patient waiting times. These distributions were gathered from hospital data and paired with actual staffing levels in order to validate their system's predicted wait times. In the model, the time required for a patient to be treated is an exogenous variable determined at the creation of the patient agent.

The Laskowski et al. model was used in the context of a single hospital to explore appropriate staffing levels for expected patient arrival rates (Laskowski and Mukhi 2009). The model, however, is also run in parallel to simulate multiple emergency rooms in order to explore policies that seek to distribute patient arrivals between hospitals in order to prevent excessive congestion in any one hospital in the system. The authors tested a policy known as random early detection, which reroutes ambulances with increasing probability as a function of the queue length of the intended destination. Reroutes are also applied to walk-in patients in the queue and can be thought of as either a suggestion that wait times might be shorter elsewhere or an offer for an ambulance to transfer them. This policy has the effect of greatly reducing congestion at any one hospital, essentially bounding the average queue length across hospitals and increasing physician utilization rate.

The Wang model focused on the total time spent in the system by patients (Wang 2009). The focus is on finding bottlenecks in the system that, if improved, could decrease total time spent in the system. The model explored the effect of adding triage nurses conditioned on the line length, then reducing staff once the line had reduced to one. The model also examined the system's sensitivity to radiology test times by having physicians order tests for a percentage of patients. The model which does not include patient acuity levels showed substantial improvements (30%) in total wait times by adding an extra triage nurse when the queue length reached ten patients. They also showed that small improvements in the average time spent in radiology (from 30 to 28 min) actually translated to much larger time savings, 10.38 min to all patients in the system.

The models described above define patient arrivals using a Poisson process, with Wang also having the option for a random uniform distribution for inter-arrival waiting times. The Jones and Evans model goes a step further in modeling patient arrival rates using a heterogeneous Poisson process with arrival rates varying each hour with daily and weekly cycles. All the models use empirical data to capture arrival rates and service times, drawing from either data available to them or data published from other ED studies. All models used between one day and one week

of simulated time as a warm up for the system to enter an equilibrium state. One of the more novel applications of agent-based models in this area was the 2009 Laskowski and Mukhi paper, in which a method for using the model to test the effects of multiple staffing levels is described. The results from each simulation are input into a genetic algorithm that generates a new set of staffing levels to be tested.

Most of the models in the papers cited above were a first attempt to replicate more complicated real-world behavior and as such included simplifications that might be relaxed in future versions of the models. The Jones and Evans model, for example, showed that when using real-world staffing levels and variable arrival rates, their model simulated with a high degree of accuracy the wait time experienced by patients. Most of the papers stated a desire to better capture the many roles of nurses in the ED. Most of the workflow models still do not incorporate much in the way of spatial components. In addition, patient outcomes (as viewed by how they leave the ED) were independent of wait time. One can imagine that instead of this being exogenous to the system, it might be a function of wait time.

2.3 Modeling Radio-Frequency IDs and Electronic Health Records

RFID is a technology that allows for both passive and active tracking of goods or patients. RFID technology is being used to help implement a real-time location system (RTLS) for patients. This system allows for patients' progress through the ED to be tracked, allowing for more accurate record keeping and potentially shifting some of the burden of such record keeping away from nurses, especially when used in conjunction with electronic health records (EHR). Such systems, however, are not without their shortcomings. Active tags, which require batteries, have a greater range than passive tags, but can cost upwards of $50 per tag. Passive tags, which do not require batteries, can cost as little as five cents and are often preferred to active tags. Passive tags typically have a limited range of approximately one to two meters, with objects and other patients potentially impeding the signal further. If, however, passive readers are placed too densely in an area, there is the potential for interference between readers.

Laskowski et al. (2010) modeled the use of RFID technology in an RTLS implemented to track patients in an ED. Given the nature of the problem, the actual topology of the ED plays a much more important role than it did in the workflow models. The topology of the ED is in two dimensions with walls, which completely block RFID signals. Patients, nurses, and other agents that represent people are represented as circles with a 60 centimeter radius. A patient is thought to be at the last RFID reader triggered until another reader is triggered. Thus, when the patient is not within a reader's radius, but has not yet triggered another reader, there is an error in where the system believes the patient is located. The paper seeks to analyze the effect of passive RFID reader placement on system accuracy. The paper, which tests

six reader configurations, runs each configuration 500 times to generate meaningful statistics (given the stochastic nature of arrival processes, this experimental design increases the likelihood that some high-volume days will be observed). There are two error terms calculated: the average Euclidean distance between a patient and the reader zone and the time a patient is reported as being in an incorrect zone. The paper then examines the effect of combining readers with variable ranges. The paper concludes that having too many readers leads to interference, which overwhelms any benefits from covering a larger area. The best performance is achieved by spacing readers as tightly as possible without overlap.

EHRs are another technology being integrated into many hospitals and emergency departments. The Poynton et al. model addresses a different question from the earlier ED workflow models, exploring how the spatial distribution of computer terminals for data entry affects workflow (Poynton et al. 2007). This model is unique in that between each stage of the patient's journey, the attending physician or nurse must enter the data from their most recent interaction with a patient. The model is clearly simplified because it assumes that caregivers are dedicated to a patient until they are discharged. The authors found that a policy based on having one computer per bed produces the best outcome. However, this strategy is often not feasible due to economic reasons; centralized computer clusters are preferable to computers distributed along the length of the hallway, depending on agent behavior.

2.4 Conclusions

ABM is an intuitive framework to design and interpret models of healthcare delivery. Agents are well suited to navigate and interact with two- and three-dimensional spatial environments, as they were considered in the RFID model. Agents have the potential to learn and adapt their behavior as the situation changes. Adaptive behavior can represent different responses to different inputs, but also different behavior to the same inputs at different times, which allows for a richer modeling of policies. The scope of applications is quite broad, ranging from the scheduling of patients, medical diagnostics, and imaging equipment to the directing of ambulances to hospitals depending on urgency, specialty, and crowding.

3 Epidemiology

The field of epidemiology is primarily concerned with how disease spreads in a given population. Studies that model the transmission of infectious diseases can focus on increasingly large scales, ranging from single hospital units or wards to entire hospitals to communities, cities, and global pandemic scales. Vector-borne diseases, or those that are transmitted by way of an intermediate carrier, are the most commonly modeled types, but airborne, waterborne, food-borne, respiratory, and

sexually transmitted diseases can also be modeled. These diseases can be modeled with different degrees of specificity to include incubation periods and periods where a colonized individual is infectious but asymptomatic. Understanding how diseases are transmitted and determining the best ways to control transmission are critical to preventing excessive spread, and epidemic modeling is an ideal method for experimenting with various strategies.

Epidemiological models have typically been compartmental models, which predict how proportions of several population states evolve over time using differential equations. The most well known of these models is the susceptible-infected-recovered (SIR) model (Kermack and McKendrick 1927), which laid the foundation for many other compartmental models (Austin and Anderson 1999; Austin et al. 1999; Beggs et al. 2008; Bootsma et al. 2006; Cooper et al. 2004; McBryde et al. 2007; Raboud et al. 2003; Robotham et al. 2007; Sebille et al. 1997). The SIR model equations are shown in (3.1), where S, I, and R represent proportions of the population that are in the susceptible, infected, and recovered states, respectively. β and γ are the transmission and recovery rates. These equations can simply be integrated over time to generate population transmission dynamics and can provide interesting results for both deterministic and stochastic scenarios. An important measure for this type of model is the basic reproduction number, R_0, which is the average number of secondary infections (i.e., transmissions) per primary case in an entirely susceptible population. R_0 is a key metric in predicting the extent to which an infection is likely to spread. If $R_0 > 1$, then an epidemic is likely to grow because on average, each infected person transmits the disease to more than one other person. If $R_0 < 1$, then an epidemic is likely to become extinct. Some models use R_0 as an input to drive transmission within a population, whereas other models use a transmission rate or probability parameter and calculate the resultant R_0 value for the population based on the number of initially and secondarily infected individuals. These models have also provided a lot of insight into the effects of certain parameters on transmission dynamics and the effectiveness of various infection control measures as defined, for example, by the following differential equations:

$$\frac{dS}{dt} = -\beta SI, \frac{dI}{dt} = \beta SI - \gamma I, \frac{dR}{dt} = \gamma I \tag{3.1}$$

Mathematical models such as the SIR model have several assumptions and limitations that prevent them from producing more valuable results. The first key assumption is that populations modeled by mathematical equations are well mixed, meaning that all individuals within the population interact with equal probability. For example, the SI term in the differential equations in (3.1) shows that transmission is proportional to the interaction between all susceptible and infected individuals, rather than a specific subset. All individuals within each compartment are assumed to be uniform, which prevents analysis of how mixed or extreme behavior, such as superspreaders or non-compliant healthcare workers (HCWs), can affect transmission dynamics. It can also be difficult to implement

time-varying or conditional behavior in a mathematical model, and thus analysis of control measures is often performed by simply varying input parameters without modeling the interactions involved in that particular intervention. Some variables, such as hand-washing probabilities and screening test return times, are reasonable to approximate with simple parameters, but others such as isolation or the effect of staffing ratios are more difficult to simplify. Several advances in mathematical modeling have been made in recent years to account for these limitations, but the results are often only valid for particular applications (Bansal et al. 2007).

Agent-based, or computational, disease spread models have expanded on the research established by the mathematical models described above and have addressed many of their limitations. They have reinforced many conclusions from mathematical and simulation models and have provided additional detail about the nature of transmission. The key advantage of ABMs is that they simulate the interactions that serve as the primary mechanism for transmission and they are capable of implementing many infection control measures explicitly. In addition, ABMs are also more adept at simulating stochastic effects, which must be captured when modeling heterogeneous populations. As a result, they have contributed significantly to a better understanding of epidemics.

The articles reviewed in this section fall mainly into two categories: process-oriented models of transmission and network models of transmission. Process-oriented models typically simulate agents that are moving through a series of stages before being removed from the population, much like they would in a DES model. Network models simulate the spread of disease between individuals using relational connections as the primary mechanism of transmission. Intermediate carriers such as HCWs do not need to be modeled explicitly in this type of model because the connections between individuals are explicit, unlike in a process-oriented model. Exemplary methods and contributions for both types of models are summarized in the following subsections.

Pandemic modeling is another research area that has been affected positively by ABM. Diseases such as malaria, SARS, smallpox, and various strains of influenza are typically modeled for these applications. Control measures are most often concerned with logistics, such as distributing vaccinations, locating community clinics, delivering emergency rations and medical supplies, and evaluating the effects of social distancing (e.g., school closures). The initial work in applying ABM techniques to pandemic scenarios was done primarily by Carley et al. (2006) and Cummings et al. (2004), which both demonstrated the value of agent-based models in generating pandemic dynamics and evaluating response strategies. Since then, agent-based pandemic models have begun to incorporate massive data sets that reflect detailed demographic, social, transportation, and even climate characteristics of a particular geographic region. However, although many of the characteristics and advantages of these models are applicable to healthcare operations management, we do not provide additional coverage of this research area because these large-scale applications are more appropriately discussed within the context of public health.
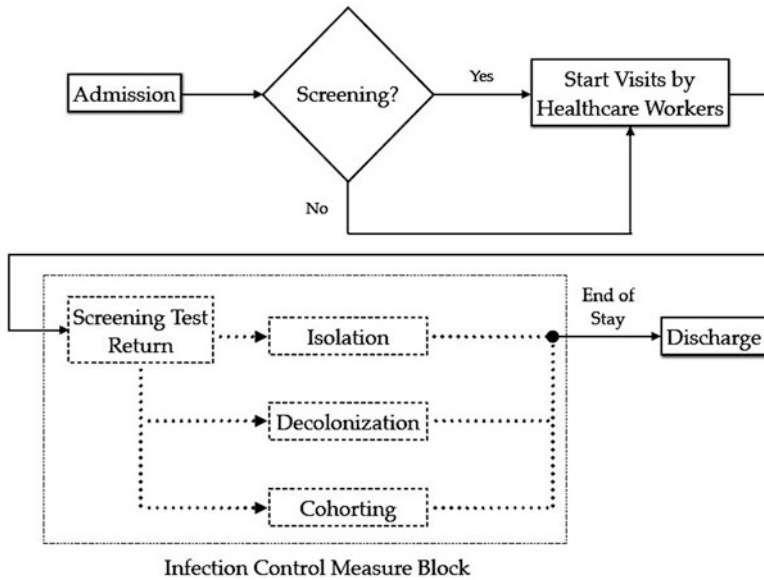
Fig. 3.2 Sample patient flow diagram of disease spread model in a hospital. Optional infection control measures are indicated by *dashed lines* and *dotted arrows*

3.1   Process-Oriented Models of Transmission

Process-oriented models of transmission are the most natural application of ABM to epidemiology. A common example is a simulation of patient-to-patient transmission in a hospital, in which patients are admitted, visited by HCWs, and discharged (see Fig. 3.2 for a sample patient flow diagram). What separates these models from traditional DES models is that there are often multiple patient and HCW types, and their behavior is often dynamic. Transmission typically occurs through HCWs, who spread an infection from one patient to another because they fail to wash their hands adequately. These types of transmissions, from an already infected (i.e., primary) patient to a newly infected (i.e., secondary) patient, are known as hospital-acquired, or nosocomial, infections. Hospital intensive care units (ICUs) are commonly modeled in an agent-based framework. ABM is ideally suited for this type of model because populations are small and diverse, patients are typically more susceptible to infection than in other hospital units, contacts between patients and HCWs are frequent and intimate, and stochastic effects are of considerable importance.

   Within agent-based models, there is often an increased ability to track various simulation data and provide additional insight into the transmission dynamics. These models can evaluate the effectiveness of various infection control measures in order to offer recommendations of which measure or bundle of measures should be implemented. Typical control measures that are modeled include the hand-washing

behavior of HCWs, diagnostic screening of patients, isolation of infected patients, vaccination, and decolonization, in which colonized patients undergo a therapeutic process that negates their ability to infect others.

There is a set of pathogens that are commonly modeled in these types of simulations, and many of these are resistant to antibiotic treatments. The most prevalent pathogens are methicillin-resistant Staphylococcus aureus (MRSA) and vancomycin-resistant enterococci (VRE), in which patients typically become colonized with the pathogen prior to developing an infection. This scenario is particularly difficult because colonized patients are often asymptomatic and, therefore, they can only become identified by using active surveillance techniques such as diagnostic screening. Consequently, these patients can spread the pathogen to HCWs and ultimately other patients before any intervention is initiated. In addition, treatment for these resistant organisms is often difficult, thus protecting patients from acquisition is the most effective approach for ensuring their safety.

The articles discussed in this section all follow a similar pattern for modeling transmission of an infectious disease in a hospital. They incorporate resistant pathogens and assess the potential effectiveness of specific control measures. However, these models differ slightly in their specific sets of experimental parameters and the nature of their results. These differences are summarized in Table 3.4.

In our review, we begin with two agent-based models that extensively explore the effects of external factors and infection control measures on disease transmission in a hospital. Both models highlight the specific level of detail afforded by ABM that is not possible using SD or DES methods. In the first example, Barnes et al. (2010) developed a detailed agent-based simulation of MRSA transmission in a hospital. The model incorporates many of the most common infection control measures and is able to reinforce the conclusions from mathematical models that one-to-one HCW-to-patient ratios and patient screening combined with isolation are most effective in reducing transmission (Austin and Anderson 1999; Austin et al. 1999; Bootsma et al. 2006; Cooper et al. 1999, 2004; McBryde et al. 2007; Raboud et al. 2003; Robotham et al. 2007). In addition, the authors highlight cases in which hospitals that have implemented infection control measures could still fall susceptible to MRSA outbreaks. Barnes et al. (2010) also conduct numerous simulations to explore the effects of individual HCW behavior on transmission, experimenting with the hand-washing probabilities of entire HCW classes (i.e., nurses and physicians) as well as rogue individuals that were less compliant than others in the same class. Many of these experiments are executed using parallel computing techniques, which alleviate the long run times caused by the model complexity. In the second example, Hotchkiss et al. (2005) also test the effects of several factors and infection control measures on transmission and demonstrate that early detection and subsequent isolation of infected patients, quick patient turnover, cohorting patients, and limiting the frequency of physician visits can all reduce the likelihood of a significant outbreak.

In addition to evaluating infection control measures and determining the most influential external factors, ABM can also provide more realistic transmission dynamics by introducing additional model complexity. This additional detail not

Table 3.4 Comparison of parameters for several significant process-oriented agent-based models of infectious disease transmission

| Model parameters	Barnes et al. (2010)	D'Agata et al. (2007)	Hotchkiss et al. (2005)	Meng et al. (2010)	Ong et al. (2008)	Temime et al. (2010)
Software platform	Python	MATLAB	Mathematica	AnyLogic	PathoSim[a]	Java
Spatial representation	Yes	No	Yes	Yes	Yes	Yes
HCW types	2	1	2	0	5	3
HCW shifts	No	Yes	Yes	No	Yes	Yes
Variable hand hygiene	Yes	*	*	No	No	Yes
Hygiene efficacy	Yes	*	*	No	No	Yes
Patient screening	Yes	No	Yes	Yes	No	No
Patient isolation	Yes	No	Yes	Yes	No	No
Decolonization	Yes	Yes	No	Yes	No	No
Variable staffing ratios	Yes	No	Yes	No	Yes	Yes
Patient cohorting	Yes	Yes	Yes	No	No	Yes
Variable transmissibility	Yes	No	Yes	Yes	Yes	Yes
Patients colonized or infected on admission	Yes	No	Yes	Yes	No	Yes
Variable # of patient visits	Yes	Yes	Yes	No	Yes	Yes
Variable patient lengths of stay	Yes	Yes	Yes	Yes	No	No
Bacterial load	No	Yes	No	No	No	No
Antibiotic resistance	No	Yes	No	No	No	No
Visitors	Yes	No	No	No	Yes	No

[a] http://www.ross-scientific.com/products.htm

*Variable hand hygiene and hygiene efficacy were not implemented explicitly, but transient HCW colonization times were variable and followed an exponential distribution

only increases the relevancy and strength of simulation results, but it can be used to address the concerns of healthcare professionals who are skeptical of the model validity. Ong et al. (2008) implemented a spatially explicit ABM of influenza transmission in a hospital unit, where at any given time, each agent occupies a specific area in the unit. Several types of HCWs are modeled, including physicians, nurses, health attendants, clerks, and cleaners. Ambulant and non-ambulant patients are also modeled. Transmission of influenza is airborne; therefore, additional model considerations must be taken into account because agents can transmit the disease without coming into direct contact with each other. Results directly related to transmission are limited, but the model generates a distribution of contacts between all pairs of agents in the unit that could be incorporated into future modeling work. The model by Meng et al. (2010) incorporates multiple routes of transmission and variable transmission rates between patients, but it produces limited results related to actual dynamics. Temime et al. (2010) describe an agent-based model of pathogenic transmission in a hospital, but the implications of these results are more appropriately addressed in the following section on network models of transmission.

Agent-based models can also be used to address questions related to antibiotic resistance, which is an important issue in disease control. Several models address the implications of antibiotic resistance at the microbiological level and are not to be discussed. However, D'Agata et al. (2007) focused on the effects of antibiotic resistance on transmission and the competition between resistant pathogens in a hospital. Single patient and HCW types are modeled explicitly, with each having eligible states of being susceptible to or colonized with resistant and/or nonresistant pathogens. The authors demonstrate that initializing decolonization treatments on patients quickly and for shorter durations can eliminate both resistant and nonresistant strains from the population. The authors are also able to develop a corresponding differential equation model that facilitates model validation and additional analysis.

3.2 Network Models of Transmission

Network models of transmission provide a different perspective than process-oriented models for analyzing the spread of infectious diseases. They are an abstract representation of the physical interactions that can lead to transmission, in contrast with the direct correlation to reality afforded by process-oriented transmission models. The edges, or connections, between nodes in the network can represent a relationship between patients in a hospital or between individuals in a community. The structure, or distribution, of these connections has a significant effect on transmission dynamics, and different network structures can be designed to represent various scenarios. Whereas process-oriented models can identify the best interventions, network models can provide insight as to where those interventions should be directed, such as targeting individuals for vaccinations or closing schools or hospital wards.
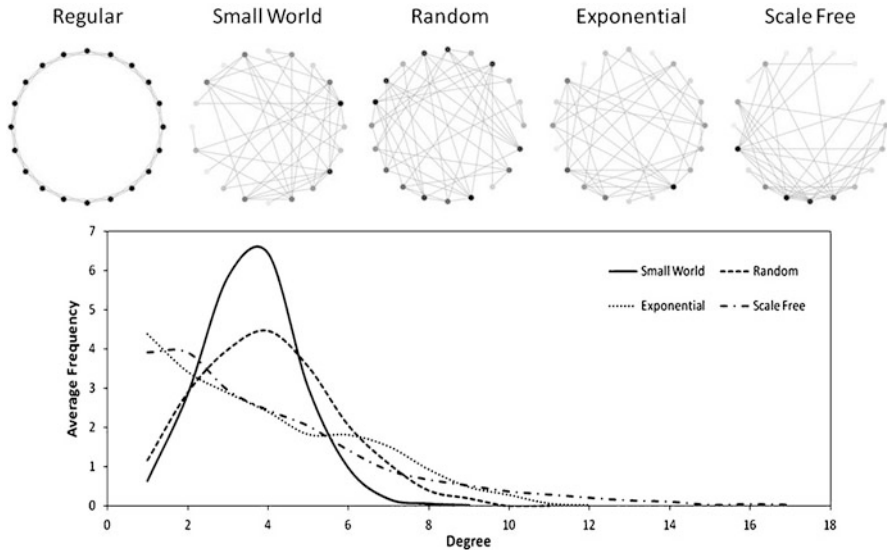
Fig. 3.3 Sample network instances for regular (degree = 4), small-world (initial degree = 4, rewiring probability $p = 0.5$), random (edge probability = 0.2105), exponential (mean degree = 4), and scale-free (generated using the Barabási–Albert model for preferential attachment) structures. *Darker shaded nodes* have a higher degree relative to the other nodes, and *lighter shaded nodes* have a relatively lower degree. The *bottom plot* shows the mean degree distribution for each network type with 20 nodes and approximately 40 edges, averaged over 100 samples. The regular network degree distribution (not shown in the graph) has a constant degree distribution, with all 20 nodes having a constant degree of 4. As the network structure changes from regular to scale-free, the degree distribution becomes more skewed, with several nodes being highly connected and the remaining nodes having relatively few connections

There are several common structures for interaction networks, which are typically characterized by the frequency distribution of node connections in the network, also known as the degree distribution (Albert and Barabási 2002). Examples of each of the following network types and their corresponding degree distributions are shown in Fig. 3.3. Regular networks have nodes that all have the same degree, and edges can be structured (e.g., nodes are connected to their nearest k neighbors) or randomly distributed to other nodes in the network. Random networks are generated by assigning an equal probability to each potential edge in the network. Each edge is then chosen at random based on the given probability, which forms a network in which the degree distribution follows a binomial model. Small-world networks (Watts and Strogatz 1998) interpolate between regular and random networks by rewiring a certain proportion of edges in a structured, regular network. Nodes in this type of network are still highly clustered, but disease could spread more quickly because there are shortcuts to other highly susceptible sections of the network. A random network is a special case of a small-world network that has rewired all of its edges. A special class of networks are exponential networks, whose degree distribution follows an exponential trend. These networks have been found

to be the most realistic structure for social interaction networks (Bansal et al. 2007). The last common type of network is a scale-free network. These networks have a power law degree distribution, in that there are a few nodes with a large number of connections and many nodes with a small number of connections. These networks are called scale-free because the average distance between any two nodes increases very slowly as the number of nodes increases. Disease transmission through these types of networks is likely to find the highly connected nodes quickly. However, transmission to the remaining population is likely to take much longer because there are fewer paths to nodes on the periphery of the network.

Mathematical models inherently assume that populations are well mixed, which means each individual has an equal probability of interacting with all other individuals. Interpretations can vary, but, in general, this assumption corresponds most closely to a regular network, whether the connections are structured or random. For certain applications, this configuration could be appropriate if each individual has approximately the same number of social contacts (e.g., child care centers). In other cases, a small-world, exponential, or scale-free network is a more accurate representation because there can be individuals that have connections to several population subgroups. These highly connected individuals often have the greatest effect on transmission, and neglecting to model them explicitly can have a significant effect on the results that are ultimately communicated to healthcare organizations.

The first set of models investigates how direct transmission can occur between individuals in a general population. These models typically consist of a homogeneous population of agents with essentially no individual characteristics other than their infection status. However, heterogeneity enters the model because the degree of each node in the network is not constant. The goal of these models is to characterize how the structure of the network affects the rate and extent of transmission, and to further demonstrate the limitations of homogeneous models.

Bansal et al. (2007) presented strong evidence that homogeneous mixing models, such as the SIR model, do not accurately predict epidemics for realistic contact networks. The authors are able to demonstrate that several empirical contact networks could all be approximated by synthetic networks with exponential degree distributions. Homogeneous mixing models, although reasonably accurate for characterizing transmission dynamics on regular, random networks, do not accurately predict epidemics on the more heterogeneous exponential and scale-free networks.

Christley (2005) focused instead on identifying the most susceptible individuals, or those most likely to become infected in the event of an outbreak, in random and small-world networks. This type of analysis is especially useful because the results could be used in developing strategies for targeting individuals for vaccination, isolation, or quarantine. The authors experiment with various measures of node centrality and determine that the degree of a node proved to be as good an indication of an individual's risk of infection as more complicated measures that would also require more information to compute. Eubank (2005) also proposed several local and global measures of network structure that could have significant implications for transmission of infectious diseases.

Fig. 3.4 The network of contacts in the modeled ICU in the Temime model. There are 18 patients and 3 types of HCWs: 2 profiles of HCWs, corresponding to nurses and physicians, are assigned to subgroups of patients, and one peripatetic-type HCW visits every patient once each day
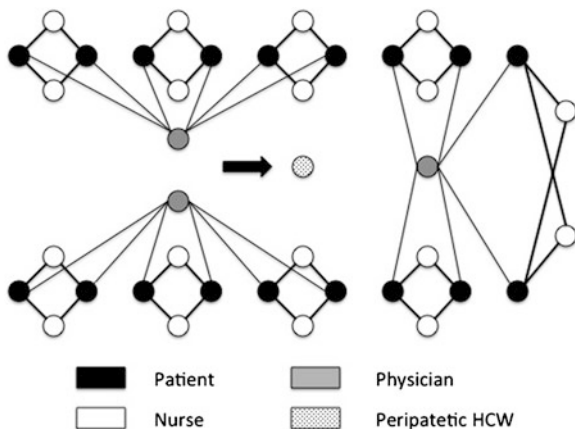
The next set of network models contain multiple types of agents that interact, whether they are explicitly or implicitly represented in the model. These studies are also focused on the structure of the network, but in addition they seek to characterize how interactions between different agent types affect transmission. The degree of heterogeneity in these network models facilitates analysis of the relative effect of each type of HCW. These models can also implement HCW behavior and bring consideration to other potential aspects of transmission such as HCW-to-HCW transmission and patient sharing. These types of interactions are not often considered, but can lead to increased levels of transmission in certain circumstances.

Temime et al. (2009) constructed a model of a hospital ICU with three types of HCWs that visit patients. Two types of HCWs are assigned to specific groups, or cohorts, of patients, whereas the third type visits all of the patients (see Fig. 3.4). The assigned HCWs represent nurses and physicians, and the third type, designated the peripatetic HCW, represents someone who could potentially come into contact with any patient, such as a nursing assistant or respiratory therapist. The model demonstrates the threat posed by the latter type and presented results that a single, noncompliant peripatetic HCW could cause the same level of transmission as if all HCWs were moderately noncompliant (i.e., 19–23% noncompliance). These effects become even more significant when HCW-to-HCW transmission occurs.

Barnes et al. (2012) also designed a network model of patient-to-patient transmission in a hospital. In this model, patients are connected directly if they share an HCW (see Fig. 3.5), which contrasts with the explicitly modeled intermediate HCW nodes in the Temime model. Nurses and physicians in this model are modeled implicitly, in that their states are only stored locally for agents in each respective cohort. By varying the number of patients, nurses, and physicians, networks are generated with different densities, calculated as the ratio of direct connections in the network to the maximum possible number of connections. Simulation experiments show that both nurses and physicians can pose threats to patients in different ways. Nurses visit patients more often, but physicians have the potential to infect patients
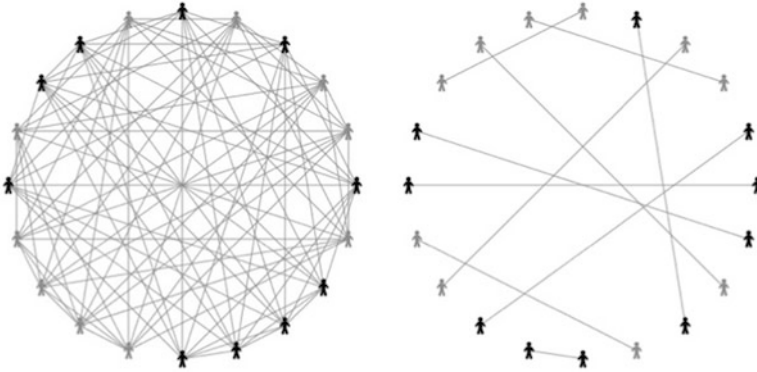
Fig. 3.5 An example of a dense (*left*) and sparse (*right*) patient network from the Barnes, Golden, and Wasil model. Patients that share a nurse are connected by a link, while patients that share a physician have the same color

in different locations in the network, similar to the threat posed by Temime's peripatetic HCW, albeit to a lesser degree. Barnes et al. also experiment with patient sharing, demonstrating that the sharing of patients between HCWs should be performed in a structured, and not random, manner.

3.3 Conclusion

Overall, ABMs are well suited for testing potential infection control measures and can provide some indication of success before implementing any particular strategy. Among ABMs, process-oriented models have a more established record of contribution, and they have produced relatively widely accepted results concerning the effectiveness of common infection control measures. The foundation has been established for network models of transmission, but there are still many areas to explore, particularly in determining which network measures are the most effective in predicting an outbreak. However, the benefit of modeling the explicit connections between individuals has been demonstrated in the unique results from this type of model. Both static and dynamic network measures may be useful, but the ease of calculation may be critical for practical use.

There are several key results that ABMs have reinforced, and several more that are unique to this methodology. The hand-hygiene compliance of HCWs is critical to preventing outbreaks of infectious diseases. In most cases, however, hand washing is not sufficient to control transmission, and additional measures such as active surveillance, patient isolation, or higher staffing ratios become necessary. Minimizing patient lengths of stay and HCW visits to patients can also reduce the risk of secondary infections, as susceptible patients are exposed less to potential infections and infected patients have fewer opportunities to infect HCWs.

ABMs have also provided insight into the relative threats posed by HCWs that care for patients in a hospital. Nurses and other HCWs that visit patients frequently are at a high risk of becoming at least transiently colonized or infected and can quickly spread disease to other patients in their care. Physicians and other HCWs that may visit many more patients than nurses can potentially create multiple pockets of infection that could lead to an entire unit becoming infected. Maintaining high staff-to-patient ratios and assigning patients to HCWs in a structured way can offset these dangers. Future studies may be able to suggest methods for identifying individuals at a high risk of infection because of their location in a network, and appropriate measures could be taken to prevent that person from becoming infected before isolation or quarantine measures become necessary.

We focused here on general models with implications for potentially many applications, but there are several models that have been applied to empirical networks (Eubank 2005; Meyers et al. 2005). These models have also achieved results that could have a significant impact, for not only their intended scenario but others. They are also good examples of the process involved in modeling real-world scenarios and can provide a good framework for future applications.

4   Healthcare Economics and Policy

In addition to modeling individual point-of-care facilities, ABMs can be used to analyze large populations at local, regional, or national levels. The interactions between different facilities are often difficult to observe from the perspective of any individual facility and are often difficult to predict. System behavior is often regulated by policies that are set in place by local, state, or federal governments. Policy changes are not usually enacted one at a time, therefore it is difficult to predict and attribute changes to a single cause when multiple changes are made. Important policy decisions should be informed using the best available data and analysis, which now includes modeling.

4.1   Healthcare Economics

In their paper, entitled "Modeling Healthcare Policy Alternatives," Ringel et al. (2010) address the benefits and limits of using agent-based (microsimulation) models to anticipate the effects of policy changes in health insurance markets. In particular, they address several issues related to the data used to populate the agents in the model and describe the relations between them and the difficulties in defining a choice model for agent decisions. The relevant models (Congressional Budget Office 2007; Girosi et al. 2009; Garret et al. 2008) seek to represent the national population in terms of households and their relationship to employers as

well as model their decisions on purchasing health insurance with respect to price, availability, and need. The scope of these models is at the national level and consists of at least three types of agents including households, employers, and insurers.

The naive approach to populating a model would be to assume independence between distributions of attributes. This assumption, however, ignores important correlations which occur in actual populations that might have a significant effect on the validity of the model, especially when dealing with socioeconomic data. Detailed information is available through the US Census and similar sources. However, simple demographic data is often not sufficient. The data sets used by each of the models contained additional information needed to provide a more detailed picture. The Survey of Income and Program Participation (SIPP), which is conducted by the US Census Bureau, was one choice used to construct a base population (Survey of Income and Program Participation SIPP). SIPP includes information on household income, taxes, and participation in government programs. Data sets such as these can also be projected onto or merged with other data sets to adjust for changes in the sample population. Other models used the Current Population Survey (CPS), which is conducted monthly by the US Census to track employment, occupation, earnings, and employee benefits (Current Population Survey CPS). The Medical Expenditure Panel Survey (MEPS), conducted by the agency for healthcare research and quality, provides data on household consumption of healthcare, including information on cost, use of healthcare, and health insurance coverage (Medical Expenditure Panel Survey MEPS). In addition to having a representative population, one must ensure that the relationships between agents are accurately described. In this case, that means an accurate relationship between firms (employers) and households (employees). This data is rarely available from the same source as the previously mentioned data. The CPS data, however, includes information on occupation. This information could be used along with a survey of firm size and industry, as well as geographic data, in order to help recreate a feasible set of firms to employ the households.

Decisions can be modeled either deterministically or probabilistically. When presented with a specific set of inputs, an agent using a deterministic decision model will always make the same decision, whereas probabilistic choice models will assign some probability to each possible choice. When given a history of agent decisions, such rules can be estimated by using a number of techniques. Decision tree learning determines a series of observations that can best recreate the choice results of the observed data. Logistic models, such as the multinomial logit model, can be used to assign probabilities to possible outcomes. The accuracy of such models is limited, however, by the amount and variety of data on previous choices. In addition, such models can break down when presented with conditions outside anything observed in the training data.

With an accurate population and sensible agent behavior, validation begins by first recreating a known set of conditions, before applying any policy changes. The Ringel et al. paper goes on to discuss potential improvements to available data and behavioral parameter estimates. These models, while primarily economic in nature,

give detailed accounts and justifications for their choices in defining their agents. These are choices that all researchers will face when attempting to build large-scale models with incomplete population and choice data.

4.2 Healthcare Policy

In addition to economic considerations, healthcare policy decisions, can have a direct or indirect impact on patient outcomes. When presented with a number of policy options, it is important to identify who will be impacted and how they will be affected. MIDAS, the Models of Infectious Disease Agent Study, is a collaborative research initiative with funding through the National Institutes of Health to model the spread of disease using agent-based models. Building on previous MIDAS models, the paper by Lee et al. (2010) uses an agent-based model to analyze the policy question of how to best prioritize, allocate, and ration vaccines during a simulated flu pandemic.

Current recommendations by the Advisory Committee on Immunization Practices (ACIP) in the USA, as stated in the paper, suggest that pregnant women, caregivers for children younger than 6 months, healthcare and emergency medical personnel, everyone between the ages of 6 months and 24 years, and people between the ages of 25 and 64 who have medical conditions associated with higher risk of complications from influenza should be prioritized to receive flu shots. These groups are listed either because they are at higher risk for complications with influenza or because they are more likely to spread influenza if infected, particularly to vulnerable populations. The paper seeks to address how strictly these guidelines should be applied in prioritizing the distribution of vaccines and whether a particular subgroup should be prioritized among the recommended groups.

The model itself is concerned with the Washington DC metropolitan area and has a population of over seven million agents, whose demographic data is drawn from the US Census Bureau's Public Use Microdata files (PUM) (Public-Use Microdata Samples PUMS). PUM files, in addition to including detailed demographic data, track transportation usage and are well suited for modeling small geographic areas. Agents interact with other agents within their household and with either schoolmates or coworkers, depending on their age. In addition, each agent is assigned a generic activity level, which varies with the day of the week. The transmission and mortality rates were calculated from previous flu pandemics. A specified percentage of infected agents require hospitalization and other medical treatments, which served as part of the basis used to calculate the total cost of the pandemic.

A series of trials using the model compared scenarios including relaxing the distribution of the vaccine to include some healthy adults between 25 and 64, prioritizing different at-risk groups, and changing the bounds of the age recommendations. The authors found that while prioritizing those what were more likely to spread the disease did reduce the overall number of infected agents, the population

also had higher total morbidity, because more at-risk patients were vulnerable to getting sick. The results as a whole support the use of the ACIP guidelines in selecting who should be vaccinated.

4.3 Conclusion

Important policy issues will continue to arise as legislatures amend existing regulations and prepare for emergency situations. It is important that detailed models be constructed to help understand the far-reaching effects that proposed changes might cause. Enacting policies such as the vaccine prioritization involves making trade-offs between the total number of infected people and the outcomes of those most vulnerable to influenza. Detailed modeling can help to illustrate the implications of the trade-offs involved with policy changes.

5 Open Opportunities for Research and Practice

ABM and simulation has been applied successfully to several applications of healthcare operations management, including healthcare delivery, epidemiology, economics, and policy. Building on the foundation established by system dynamics and DES, ABM has provided new insight to these problems by modeling individuals and the interactions between them. This perspective has facilitated analysis at both the individual and system levels, which is not typically possible using other methods. The greatest value of ABMs is that they can be used as virtual environments to evaluate policy alternatives, some of which would be infeasible or unethical to experiment with in practice. This capability can help healthcare organizations make better decisions, which can potentially lead to a higher quality of care for patients and extensive cost savings.

As with the other modeling methodologies, the challenge in constructing ABMs is first generating realistic dynamics for a particular scenario and then analyzing those patterns to identify the best strategies for intervention. Determining the appropriate level of detail for an agent-based model can be a difficult challenge. Extremely detailed models require many parameters, and it is difficult to get accurate values for these parameters from empirical data and the research literature. Parameters are often set at reasonable values, and sensitivity analysis is performed to address any shortcomings associated with poor selection, but this type of analysis can become very time-consuming if model run times are large. Models with few parameters mostly avoid this problem, but are sometimes challenged due to their lack of detail and complexity. One solution to this issue is to use hybrid models, which incorporate system dynamics and agent-based methods. In a hybrid model, one might model a network of agents whose states are affected by both their interaction with other agents and an internal differential equations

model. These models can take advantage of the strengths of each methodology, most importantly the lower computational demands of system dynamics and the heterogeneity afforded by ABM. For all levels of fidelity, agent-based models can incorporate and generate large quantities of data. As a consequence, statistical rigor, efficient data analysis techniques, and visualization are all critical to producing insightful results and communicating those findings to healthcare professionals.

Several models focused on only a particular hospital configuration or network instance, which may limit the usefulness of a particular strategy. Certain trends may persist over similar variations, but the strongest impact is likely to be derived from models that suggest robust solutions that have implications for many scenarios. Other models used parameterized configurations, which are more broadly applicable, but may not correspond well to any particular situation. ABMs can also operate on a variety of time scales, and it is important to choose an appropriate scale that is relevant to the problem and compatible with model parameters.

Another challenge with ABMs is related to validation. The level of validation required for a particular agent-based model is primarily driven by its ultimate purpose. Models that are intended to produce accurate quantitative (i.e., predictive) results may require extensive validation, whereas more qualitative (i.e., illustrative) models have less stringent requirements. Validation is difficult for some applications because they lack empirical data as a baseline for comparison. Therefore, it is not always possible to know what simulation outcomes should look like. When empirical data are available, there are several approaches for using them to validate a model (Fagiolo et al. 2007). The best of these approaches primarily involves incorporating empirical data, historical perspective, and subject matter expertise in the selection of the best possible set of model assumptions, parameter settings, and initial conditions. By combining these data sources, agent-based models are more likely to use appropriate values for input parameters and generate reasonable system responses. For any validation approach, it is always necessary to perform sensitivity analyses to identify any parameters that cause nonproportional changes in the system.

Empirical data is typically available for many applications within healthcare operations, especially those related to patient information and resource management. In addition, subject matter expertise is readily available, which should be leveraged to develop models with an appropriate set of agent characteristics and behaviors. However, validation becomes more difficult when the underlying dynamics are not well understood or there is no historical reference, as is generally the case for the transmission of antibiotic-resistant pathogens or the implementation of a novel healthcare policy. For these cases, an agent-based model may have more value in demonstrating relative trends or generating alternative outcomes for comparison, and less value as a predictive model.

Moving forward, new and existing agent-based models should continue to build on the most relevant achievements of other models, including those from other fields that have taken advantage of the methodology. Making these models widely available is a practice that can help to accelerate progress. This evolutionary process is another advantage of ABM, as system dynamics and DES models are more

difficult to augment with additional complexity. Each individual model may have limitations, but as the research grows, many of those limitations can eventually be eliminated, thereby increasing the acceptance of these methods by healthcare organizations.

References

Albert R, Barabási AL (2002) Statistical mechanics of complex networks. Rev Mod Phys 74(1): 47–97

Austin DJ, Anderson RM (1999) Studies of antibiotic resistance within the patient, hospitals and the community using simple mathematical models. Phil Trans Roy Soc Lond B 354(1384): 721–738

Austin DJ, Bonten MJ, Weinstein RA et al (1999) Vancomycin-resistant enterococci in intensive-care hospital settings: Transmission dynamics, persistence, and the impact of infection control programs. Proc Natl Acad Sci USA 96(12):6908–6913

Bansal S, Grenfell BT, Meyers LA (2007) When individual behaviour matters: Homogeneous and network models in epidemiology. J Roy Soc Interface 4(16):879–891

Barnes S, Golden B, Wasil E (2010) MRSA transmission reduction using agent-based modeling and simulation. INFORMS J Comput 22(4):635–646

Barnes S, Golden B, Wasil E (2012) Exploring the effects of network structure and healthcare worker behavior on the transmission of hospital-acquired infections. IIE Tran Healthc Syst Eng 2:259–273

Beggs CB, Shepherd SJ, Kerr KG (2008) Increasing the frequency of hand washing by healthcare workers does not lead to commensurate reductions in staphylococcal infection in a hospital ward. BMC Infect Dis 11:1–11

Bootsma MCJ, Diekmann O, Bonten MJM (2006) Controlling methicillin-resistant Staphylococcus aureus: Quantifying the effects of interventions and rapid diagnostic testing. Proc Natl Acad Sci USA 103(14):5620–5625

Carley KM, Fridsma D, Casman E, Yahja A, Altman N, Chen L-C, Kaminsky B, Nave D (2006) BioWar: Scalable agent-based model of bioattacks. IEEE Trans Syst Man Cybern A 36(2): 252–265

Christley RM (2005) Infection in social networks: Using network analysis to identify high-risk individuals. Am J Epidemiol 162(10):1024–1031

Congressional Budget Office (2007) CBO's Health insurance simulation model: A technical description. www.cbo.gov/ftpdocs/87xx/doc8712/10-31-HealthInsurModel.pdf. Accessed on August 2011

Cooper BS, Medley GF, Scott GM (1999) Preliminary analysis of the transmission dynamics of nosocomial infections: Stochastic and management effects. J Hosp Infect 43(2):131–147

Cooper BS, Medley GF, Stone SP et al (2004) Methicillin-resistant Staphylococcus aureus in hospitals and the community: Stealth dynamics and control catastrophes. Proc Natl Acad Sci USA 101(27):10223–10228

Cummings D, Burke DS, Epstein JM, Singa RM, Chakravarty S (2004) Toward a containment strategy for smallpox bioterror: An individual-based computational approach. Brookings Inst Pr, Washington, DC, pp 1–55

Current Population Survey (CPS). www.census.gov/cps/. Accessed on August 2011

D'Agata EMC, Magal P, Olivier D et al (2007) Modeling antibiotic resistance in hospitals: The impact of minimizing treatment duration. J Theor Biol 249(3):487–499

Eubank S (2005) Network based models of infectious disease spread. Jpn J Infect Dis 58(6): S9–S13

Fagiolo G, Moneta A, Windrum P (2007) A critical guide to empirical validation of agent-based models in economics: Methodologies, procedures, and open problems. Comput Econ 30(3):195–226

Fone D, Hollinghurst S, Temple M et al (2003) Systematic review of the use and value of computer simulation modelling in population health and health care delivery. J Publ Health 25(4):325–335

Garret B, Clemans-Cope L, Bovbjerg R, Masi P (2008) The Urban institute's microsimulation model. www.urban.org/url.cfm?ID=411690. Accessed on August 2011

Girosi F, Cordova A, Eibner C, Gresenz C, Keeler E, Ringel J, Sullivan J, Bertko J, Buntin M, Vardavas R (2009) Overview of the COMPARE microsimulation model. www.rand.org/pubs/working_papers/WR650. Accessed on August 2011

Hotchkiss JR, Strike DG, Simonson DA et al (2005) An agent-based and spatially explicit model of pathogen dissemination in the intensive care unit. Crit Care Med 33(1):168–176

Jones S, Evans R (2008) An agent based simulation tool for scheduling emergency department physicians. AMIA Annu Symp Proc 2008:338–342. Published online 2008. PMCID: PMC2656074

Jun JB, Jacobson SH, Swisher JR (1999) Application of discrete-event simulation in health care clinics: A survey. J Oper Res Soc 50(2):109–123

Kanagarajah A, Lindsay P, Miller A, Parker D (2006) An exploration into the uses of agent based modeling to improve quality of health care. In: Minai A, Braha D, Bar-Yam Y (eds) Proceedings of the 6th international conference on complex systems, Boston, MA

Kermack WO, McKendrick AG (1927) A contribution to the mathematical theory of epidemics. Proc Roy Soc Lond A 115:700–721

Koelling P, Schwandt MJ (2005) Health systems: A dynamics system benefits from system dynamics. In: Kuhl ME, Steiger NM, Armstrong FB, Joines JA (eds) Proceedings of the 2005 Winter Simulation Conference, Orlando, FL, USA, pp 1321–1327

Laskowski M, Mukhi S (2009) Agent-based simulation of emergency departments with patient diversion. Electronic healthcare. Springer, Berlin, pp 25–37. Isbn: 978-3-642-00413-1

Laskowski M, McLeod RD, Friesen MR, Podaima BW, Alfa AS (2009) Models of emergency departments for reducing patient waiting times. PLoS ONE 4(7):e6127. doi:10.1371/journal.pone.0006127

Laskowski M, Demianyki B et al (2010) Uncertainties in RFID tracking systems in an emergency department. Health Care Management (WHCM), 2010 IEEE Workshop

Lee B, Brown S, Korch G, Cooley P, Zimmerman R, Wheaton W, Zimmer S, Grefenstette J, Bailey R, Assi T, Burke D (2010) A computer simulation of vaccine prioritization, allocation, and rationing during the 2009 H1N1 influenza pandemic. Vaccine 28(31):4875–4879

Macal C, North M (2007a) Managing business complexity: Discovering strategic solutions with agent-based modeling and simulation, 1st edn. Oxford University Press, New York

Macal CM, North MJ (2007b) Agent-based modeling and simulation: Desktop ABMS. In: Henderson SG, Biller B, Hsieh M-H, Shortle J, Tew JD, Barton RR (eds) INFORMS Winter Simulation Conference, Washington, DC, USA, pp 95–106

McBryde ES, Pettitt AN, McElwain DLS (2007) A stochastic mathematical model of methicillin resistant Staphylococcus aureus transmission in an intensive care unit: Predicting the impact of interventions. J Theor Biol 245(3):470–481

Medical Expenditure Panel Survey (MEPS). www.meps.ahrq.gov/mepsweb/. Accessed on August 2011

Meng Y, Davies R, Hardy K, Hawkey P (2010) An application of agent-based simulation to the management of hospital-acquired infection. J Simul 4(1):60–67

Meyers LA, Pourbohloul B, Newman MEJ, Skowronski DM, Brunham RC (2005) Network theory and SARS: Predicting outbreak diversity. J Theor Biol 232(1):71–81

Ong B, Chen M, Lee V, Tay J (2008) An individual-based model of influenza in nosocomial environments. In: Bubak M, van Albada GD, Dongarra J, Sloot PMA (eds) Comput Sci, Int Conf Comput Sci 2008, Part I. Lecture Notes in Comp Sci 5101:590–599

Pearce B, Huynh N, Harris S (2010) Modeling interruptions and patient flow in a preoperative hospital environment. Proceedings of the 2010 Winter Simulation Conference Baltimore, MD, USA

Poynton M, Shah V, BeLue R, Mazzotta B, Beil H, Habibullah S (2007) Computer terminal placement and workflow in an emergency department: An agent-based model. Proceedings of the complex systems summer school, Winter Simulation Conference Washington, DC, USA

Public-Use Microdata Samples (PUMS). www.census.gov/main/www/pums.html. Accessed on August 2011

Raboud J, Saskin R, Simor A et al (2003) Modeling transmission of methicillin resistant Staphylococcus aureus among patients admitted to a hospital. Infect Contr Hosp Epidemiol 26(7):607–615

Ringel J, Eibner C, Girosi F, Cordova A, McGlynn E (2010) Modeling health care policy alternatives. Health Serv Res 45:1541–1558

Robotham JV, Jenkins DR, Medley GF (2007) Screening strategies in surveillance and control of methicillin-resistant Staphylococcus aureus (MRSA). Epidemiol Infect 135(2):328–342

Sebille V, Chevret S, Valleron JA (1997) Modeling the spread of resistant nosocomial infections in an intensive-care unit. Infect Contr Hosp Epdemiol 18(2):84–92.

Survey of Income and Program Participation (SIPP). www.census.gov/sipp/. Accessed on August 2011

Temime L, Opatowski L, Pannet Y et al (2009) Peripatetic health-care workers as potential superspreaders. Proc Natl Acad Sci USA 106(43):18420-5

Temime L, Kardas-Sloma L, Opatowski L et al (2010) NosoSim: An agent-based model of nosocomial pathogens circulation in hospitals. Proc Comput Sci 1(1):2245–2252

Wang L (2009) An agent-based simulation for workflow in emergency department. In: Proceedings of the 2009 IEEE systems and information engineering design symposium, University of Virginia, Charlottesville, VA

Watts DJ, Strogatz SH (1998) Collective dynamics of small-world networks. Nat 393(6684): 440–442