The impact of the residency teaching model on the efficiency of the emergency department at an academic center

Q4 John Silberholz [a], David Anderson [a,*], Bruce Golden [a,1], Michael Harrington [b], Jon Mark Hirshon [c]

[a] R.H. Smith School of Business, University of Maryland, College Park, MD 20742, USA
[b] University of Maryland Medical Center, Baltimore, MD 21201, USA
[c] University of Maryland School of Medicine, Baltimore, MD 21201, USA

ARTICLE INFO

ABSTRACT

The residency teaching model is often cited as a source of inefficiency in the healthcare system. We build a simulation model of an Emergency Department (ED) at a large urban academic hospital. Using historical data and a natural experiment involving residents in the ED, we show that residents in fact increase throughput and lower service and waiting times compared to not being there at all.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

The rising cost of healthcare is of significant social and political concern in America today. According to the Center for Medicare and Medicaid services (CMS), in 2007, total healthcare spending in the USA was $2.2 trillion. At $7421 per person, this accounts for over 16% of the nation's GDP. By 2007, healthcare spending was more than three times what it had been in 1990 [16]. CMS expects that healthcare spending will keep increasing at 6% per year until 2018, at which point it will account for a fifth of the United States economy. While there are many components to these costs, hospitals contributed the largest amount, at 32% of all healthcare expenditures. Increasing hospital efficiency is one way to help slow the growth of healthcare spending [7]. reports that total healthcare spending increased to $2.5 trillion or $8047 per person, by 2009.

One potential source of inefficiency that we will study in this paper is the residency teaching model. After students graduate from medical school, they must complete three to seven years of additional training under a senior doctor, called an attending physician, to become board certified in a medical or surgical specialty. Residencies can be completed in any general or specialty field within medicine or surgery. Upon successful completion of residency and the specific medical boards for that specialty, a doctor is then considered a certified specialist. This level of training is required to work as an attending physician in an academic center. While working as a resident, a new doctor will diagnose and treat patients under the supervision of an attending physician who oversees and teaches the residents, while providing clinical care. It is a common hypothesis that the presence of residents in a hospital setting hinders overall system efficiency [5]. Because attending physicians have to spend time teaching residents that could be spent treating patients directly, it has been suggested that residents slow down treatment and hinder efficiency.

We worked with the University of Maryland Medical Center (UMMC) to help determine the impact of their residency teaching model on efficiency in the Emergency Department (ED). We collected data and designed and implemented a simulation model of the ED. In Section 2, we review the relevant literature. In Section 3, we discuss our data and provide a detailed description of the simulation model. Validation of the model is given in Section 4. In Section 5, we discuss the results and implications. The conclusions are presented in Section 6.

2. Literature review

In this section, we discuss studies about the effects of residents on ED efficiency. The resident education model creates a dual role for attending physicians in the ED, because a resident's role includes both treating patients and learning medicine. Thus, the resident care model can affect patient throughput because of the additional time spent on instruction.

Recent research has found that residents do decrease efficiency in hospital settings. In one study, researchers aimed to review ED

Q1 * Corresponding author. Tel.: +1 7248774139.
E-mail addresses: davidryberganderson@gmail.com (D. Anderson), bgolden@rhsmith.umd.edu (B. Golden).
¹ Tel.: +1 301 405 2232; fax: +1 301 405 8655.

patient waiting times, time until an admission decision was made, and total ED length of stay during periods when residents were on strike versus times of normal resident staffing patterns [12]. They found that without residents, the ED had higher throughput and the length of stay was reduced [17]. examined the effect of adding residents to an ED at a community hospital. They found that there was a weak, positive correlation between ED patient length of stay and the presence of residents [4]. studied the efficiency of residents as they gained experience. They found that as residents become more experienced they increase their throughput [21]. observed the effects of a resident strike on quality and throughput in an ED at a large teaching hospital. They found that replacing residents with staff physicians led to an increase in throughput and in quality of care.

Other studies, however, have shown that residents have no negative effects on throughput or treatment times [5]. authors looked at the introduction of anesthesiology residents to surgical wards and expected to find decreased efficiency. However, they found no significant adverse economic or health effects [20]. studied the addition of residents to a trauma care center and concluded that residents improved efficiency while having no effect on the quality of care.

Methodologically, our paper relies on simulation modeling and queueing theory. These methods have been used extensively in the hospital operations management literature [6,14,15].; and [1] provide surveys of simulation models used in healthcare research. Simulation has a wide variety of applications in healthcare, such as modeling patient flow [2], optimizing resource allocation [18], and evaluating surgery scheduling strategies [3].

Queueing theory is another technique widely used in the hospital operations management literature [10]. and [8] provide surveys of applications of queueing theory to healthcare problems. For example, queueing theory has been used in the emergency department to determine appropriate staffing levels in order to reduce the proportion of patients who leave without being seen [11] and to assist in bed management planning [9].

3. Data and simulation model

We were motivated by the inconclusive literature to study whether residents help or hurt efficiency in the ED. At the UMMC, every Wednesday morning there was a seminar that the residents had to attend, so they were not present in the ED. Because of this, patients who were treated on Wednesday mornings were not seen by a resident, but only by attending physicians. This observation (residents present vs. not) suggests a natural experiment to determine what effect removing residents would have during other parts of the week. We designed a simulation model to exploit this natural experiment. This research was reviewed and approved by the Institutional Review Board at the University of Maryland, Baltimore.

Because there are no changes in staffing levels in the ED on Wednesday mornings other than the presence or absence of residents, the differences in treatment times for similar patients can be attributed entirely to the presence or absence of residents. Typically, there are two attending physicians on duty and four or five residents in the ED. When the residents are present, they do almost all of the "hands-on" treatment of patients, while the attending physicians play a managerial/supervisory role. When the residents are present, they are simultaneously treating patients and receiving instruction from the attending physicians. The attending physicians oversee the care and teach the residents. Therefore, our simulation model assumes that when residents are present they treat every patient who arrives. When the residents are absent, due to the seminar, the attending physician's role shifts from supervisory to active care-providing. As a consequence, they now spend their time treating patients, rather than supervising and teaching residents. The changes in treatment times that we see when residents are not present are a result of this shift. This assumption was motivated by conversations with ED physicians at the UMMC.

To attribute treatment time changes on Wednesday mornings to staffing levels, we must verify that Wednesday mornings are similar to the rest of the week in terms of arrival rates and patient severity. To do this, we compare the patients who arrive on Wednesday mornings (when residents are absent) to those who arrive at all other times of the week (when residents are present). Fig. 1 shows the historical arrival rates over the course of the week. There is a wide range of arrival rates for Wednesday mornings. In general, there are more arrivals than on weekend mornings and fewer than on Monday or Tuesday mornings. In addition, morning arrival rates are higher than overnight rates and lower than afternoon rates. So, Wednesday morning arrival rates are not atypical in any way. Furthermore, the patient population mirrors that of the rest of the week, in terms of severity and admission rate. We compared the two patient populations (those treated when residents were present and those treated when they were absent). We found that when residents were absent, 47%, 50%, and 3% of patients were of high, medium, and low severity, respectively, while those numbers were 45%, 51%, and 4% when residents were present. A Chi−Square test fails to reject ($p = .81$) the hypothesis
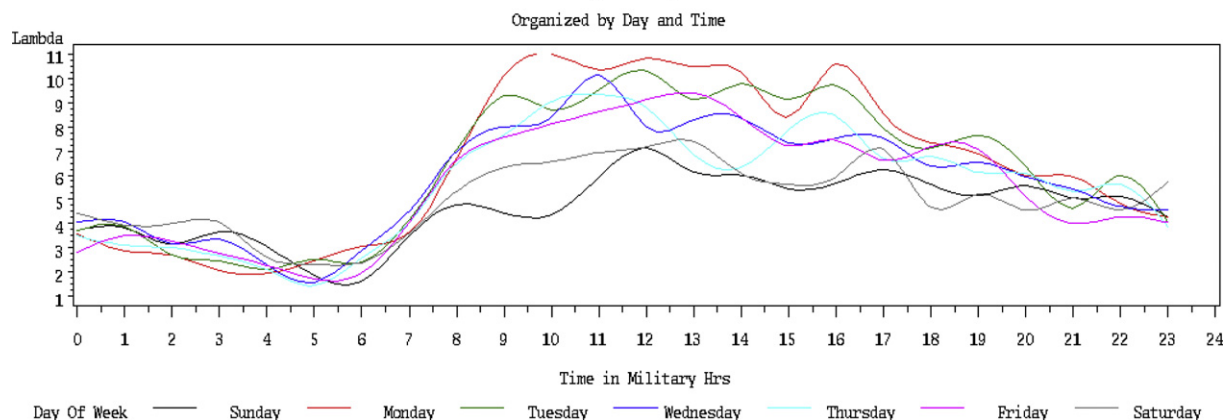

Fig. 1. Arrival rates by day of week.

that the underlying severity distribution is the same between the two patient populations. Similarly, the proportion of patients needing lab tests (72.4% vs. 75.8%) is not provably different ($p = .22$) between the two populations. The fact that the two patient populations are so similar gives us further confidence that the differences that we observe in treatment times between the two populations is caused by the presence or absence of residents, and not by other factors.

We are, therefore, fortunate to have a representative sample of patients not treated by residents on Wednesday mornings, which enables us to quantify the effect of having residents work in the ED.

Based on historical arrival and severity data, we built a simulation model of the ED. Fig. 2 shows a flow diagram of the ED simulation. We use this model to determine the effect of residents not just on treatment times for patients, but on the ED system as a whole. By building a simulation model, we can show how the presence of residents in the ED affects waiting times, throughput, and total time in the system. Moreover, the ED is a complex system with many interdependent parts. Because of this complexity, we felt that a simulation model would be more appropriate than other types of models (e.g., queueing models). Building a simulation model also allows us to easily experiment with the system to see how changing parameters of the system would affect performance. We implemented the simulation using SimPy, a discrete-event simulation language for Python.

The effect that residents have on treatment times is handled implicitly by the simulation. As discussed previously, we assume that every patient treated when residents are present is treated by a resident, while those treated when residents are absent are not. We do not model the specific movements of individual physicians through the ED or every doctor–patient interaction. Instead, we take a higher-level view of the ED and simply simulate patient flow.

For this study, we used historical data from the UMMC ED. The UMMC ED is divided into separate sections that treat adult medical patients, pediatric patients and psychiatric patients. There is a separate area outside the ED for patients with significant trauma. The main adult ED, the site for the prospective data collection, sees approximately 50,000 primarily adult medical and urgent care patients annually.

We build our model from the UMMC patient database data from October 1, 2009 to January 31, 2010 that contained data from the adult medical and psychiatric areas. The patient identities were masked. There were almost 17,000 patient visits during these four months and each record contained information about the patient's triage score, treatment process, and when and why they left the ED.

3.1. Patient creation

Patients enter the simulation model according to a nonhomogeneous Poisson process, with the arrival rate based both on time of day and day of the week, drawn from the historical arrival data. After each patient is generated, he/she is seen by the triage nurse. At the triage station, the patient is assigned a severity score from 1 (highest) to 5 (lowest), and held for a random amount of time based on historical average triage times. A small number of patients are not given a severity score. These correspond to patients brought in via ambulance and with extremely high severity. In addition to the severity score, the simulation determines the amount of lab work the patient needs and whether or not the patient will eventually be admitted to an inpatient ward, based on the severity score.

We chose these three attributes (severity, labs, and admission) because they were the most important in determining the treatment time that a patient required and the most medically relevant. Higher severity patients take, on average, longer to treat. A high-severity patient will require more intensive care and will be held longer in the treatment bed. Similarly, a patient who is admitted to the hospital from the ED is likely to be held longer. Patients who are admitted have more severe and complex problems than those who are not. Lastly, the number of labs that a patient needs directly affects the treatment time. Lab work takes time to process, which causes the patient to stay longer.

3.2. Patient selection

Once a patient is discharged and a bed becomes free, the physician must select a patient from the waiting room. While we might expect the patients to be selected strictly according to severity, the historical data confirms that this is not the case. Based on the historical data, we found that the number of times that a patient was passed over lowered his future chances of being selected for treatment. This means that a severity 2 patient who has been passed over a few times might be less likely to be picked than a newly arrived severity 3 patient, even though he is in a higher severity class.

There is no deterministic rule for how patients are selected, so we constructed a discrete choice model, using logistic regression, to model how patients were selected. Patients were split into 4 severity categories: high (severity score of 1 or 2), medium (score of 3), low (score of 4 or 5) and N/A (no score given). Within each of the severity categories, we split the patients again 4 ways, based on how many times they had been passed over in the selection process: never, once, 2–3 times, and 4 + times, giving us 16 different patient categories. The probability that each patient would be chosen from a waiting room with one patient of each type is shown in Fig. 3. We see that high severity patients are much more likely to be chosen than low severity patients, but also the more times a patient has been passed over the less likely he/she is to be selected.

This presented us with a discrete choice problem. Each time a bed becomes free, triage nurse must select one and only one patient from the waiting room. Each time a patient is selected, in

Fig. 2. Flow diagram of simulation.

Patients Generated → Triage → Waiting Room → Treatment → Discharge

Fig. 3. Probability each patient type is chosen from the waiting room.

the historical data, we note which patient type is selected, and how many of each type are still in the waiting room. Typically this class of problems is solved using multinomial logistic regression. However, this approach requires on the order of 2^N terms in a choice set with N alternatives. In our case, this would require estimating more than 65,000 terms, which is computationally prohibitive. Instead, from the constructed dataset, we built a series of logistic regression models (see [13] that measure the probability of each type of patient being selected given the distribution of patients in the waiting room. The probabilities from these regressions were used to choose which type of patient would be selected next in the simulation model. These sequential logistic regression models approximate what multinomial logistic regression does.

Because patients sometimes leave the waiting room before being treated, our simulation must take abandonment into account. From the historical data, we know the probability that a patient of a given severity will still be in the waiting room based on the number of hours he/she has been waiting. After a patient is selected from the waiting room to be treated, we determine if he/she is still in the waiting room. If the patient is absent, another patient is selected from the remaining patients in the waiting room. Once a patient has been selected and is still present, he/she is assigned to a treatment bed and held until treatment is over. The probability that a patient of each severity class is still in the waiting room is plotted in Fig. 4. The curves are not smoothly decreasing because the sample size becomes very small as waiting times increase. Very few patients wait over 6 h to be seen, and the data only record whether or not the patient was present when selected, not the exact time that they left the waiting room. We only know when patients who have left without being seen are called to be placed in a bed.

3.3. Treatment time

Once in the treatment bed, the patient remains there for a length of time drawn from empirical distributions. We used empirical distributions because they were able to capture the long

Fig. 4. Percent remaining vs. time until called back.

tails of treatment times better than kernel densities. When possible, we categorized each patient by a number of binary splits. The first split was based on whether or not the ED was congested (defined as more than 4 patients in the waiting room). Second, we split the patients based on whether or not they were eventually admitted to the inpatient ward, as admitted patients and discharged patients have different ED length of stays and different service needs. Lastly, we split the patients based on the amount of lab work they needed, their severity level, and whether or not they were seen by a resident. However, due to data sparseness issues, we weren't able to make every split. For example, there were very few low severity patients with no lab tests who were admitted to the inpatient ward. The length of treatment time for each patient was drawn from the empirical distribution for that patient's category.

Once the treatment time had elapsed, the patient left the simulation (either via discharge or admittance to the inpatient ward), and the bed was held shortly while being prepared for a new patient. Once the bed has been cleaned, a new patient is called back, and the cycle repeats. Because one of the parameters that determine treatment time is whether or not the patient is treated by a resident, we can run experiments with our simulation by varying that parameter for groups of patients. In Section 5, we present these experiments.

By measuring treatment times based on the treatment and ED characteristics (labs, severity, congestion, admission to the hospital), we are able to control for possible confounding of the effect of residents, enabling us to isolate the effect that residents have on ED efficiency. For instance, whenever a simulated low-severity patient with no lab tests enters the ED during an uncongested time with residents present and is later discharged, we draw treatment times for that patient from an empirical distribution of all similar patients in the historical ED data who were treated when the ED was uncongested and when residents were present (all times except Wednesday mornings). If we were simulating the same patient, except without residents present, we would draw treatment times from an empirical distribution of all similar patients who were treated when the ED was uncongested and when residents were not present (Wednesday mornings).

4. Validation

After building the simulation model, we tested it to make sure that it was a valid replication of the system we were simulating. We did this by comparing the similarity of the outputted data from our model with the observed performance of the ED. While validating the model, we mirror the actual system, with residents present all week except for Wednesday mornings. We compared statistics from the simulation regarding patients per bed per day, the rate at which patients abandoned the waiting room before being seen, time spent until placed in a bed, and total time in system with those from the historical database. These are metrics often used to evaluate ED performance and efficiency. By demonstrating that the data generated for these metrics were statistically similar to the data from the historical database, we were able to confirm that we have a valid simulation model.

We chose the above-mentioned comparison metrics because they describe the overall performance of the ED. We simulated 20 years of data to compare to the historical values. From the simulated data, we calculated the mean and standard deviation for each of the performance metrics. Table 1 illustrates the similarities between the simulation model and the historical data. None of the metrics were provably different from the historical values.

We used a Kolmogorov–Smirnov (K–S) test to test the similarity of the total time in system distributions from the simulation and
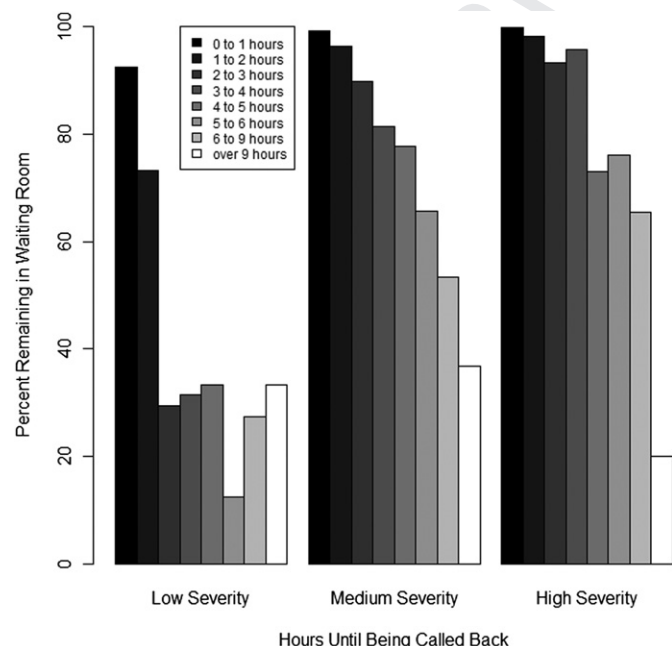
Table 1
Comparison of simulated means with historical means for key ED efficiency metrics.

Metric	Historical mean	Simulation mean	P-value
Patients per bed per day	2.38	2.39	0.4413
Abandonment rate (in percent)	8.02	7.76	0.4611
Time to first bed placement (in minutes)	80.32	81.12	0.8650
Total time in system (in minutes)	550.15	549.28	0.9134

the historical data. The K–S statistic for two samples measures the difference between the empirical cumulative distribution functions (ECDFs) of the two samples. The ECDFs are step functions that approximate the underlying distributions from which the samples are drawn. We find the maximum vertical distance between the two ECDF curves, and compare it to the expected difference if the two samples were drawn from the same population. If they are farther apart than what would happen in 95% of cases, we can say with 95% confidence that the two samples were drawn from different distributions.

The K–S test statistic for the total time in system metric was .0075, meaning that the farthest distance between the two distributions was 0.75%. This translates to a *p*-value of .513, meaning that we cannot reject the null hypothesis that the simulation output and the historical data have the same length of stay distribution. Our time in system distribution matches the historical data almost perfectly, and the other performance metrics are similar to the historical data at the means. Noting that simulations by their very nature simplify a complex system and, therefore, cannot perfectly replicate that system's performance, we felt comfortable with the model validation results.

5. Experiments and results

In our first experiment, to determine the effect that residents have on ED efficiency, we varied the proportion of patients seen by residents from 0 to 1, in increments of 0.1, and observed the changes in efficiency metrics, such as throughput and average waiting time. From one run to the next, the only change in the system is the percent of patients seen by a resident. In this experiment, residents see each patient with the same probability, regardless of patient severity. Because treatment by a resident is a parameter in the simulation, we randomly select whether a patient is treated by a resident when that patient enters the ED. We ran 20 years' worth of simulations for each level of resident presence and recorded the performance metrics from these simulations. These experiments test the hypothesis that the addition of
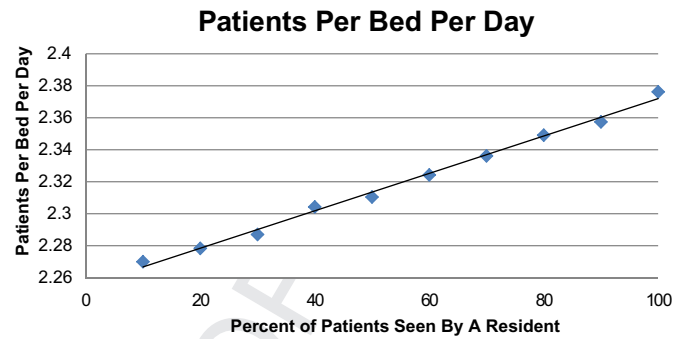
**Fig. 6.** Total throughput (in patients per bed per day) vs. resident presence.

residents to the ED slows down doctor performance and harms system efficiency.

We found strong linear trends in the relationship between the patient-based characteristics and the presence of residents. For example, we saw decreases of over 16% in total time (from 11.5 to 9.5 h) for both high- and low-severity patients when residents were added. Additionally, we saw decreases in the time to get patients into a bed of 23% for high-severity patients and 20% for low-severity patients. Fig. 5 shows the relationship between the total time in the ED for patients and the percent of patients treated by residents.

We also observed increases in system-wide efficiency. We measured total throughput in terms of patients treated per bed per day and found that having residents treat patients helped improve throughput. In particular, we found a 6% increase in total throughput (from 2.26 patients per bed per day to 2.38) when resident presence was increased from 0 to 100%. Fig. 6 shows a plot of patient throughput versus resident presence.

The third performance metric we monitored was time to first bed. Again, we found that increasing the fraction of patients seen by a resident helped to improve system performance. This is especially important in an ED because patient welfare often depends on how quickly they can be seen and treated by a doctor. Fig. 7 shows the effect of increasing the percent of patients seen by a resident on time to first bed. The addition of residents lowers average waiting times by 35% (from 92 min to 60).

In our second experiment, we independently varied both the percentage of high- and low-severity patients seen by residents. Because the residents' main purpose in the ED is to learn, and because the high-severity cases are the most instructive, we fixed the fraction of high severity patients seen by residents always at or above one half. We simulated 20 years with 121 different patient mixes, varying the fraction of high severity patients seen from .5 to

Fig. 5. Total time in system (in hours) vs. resident presence.

Fig. 7. Average time to first bed (in minutes) vs. resident presence.

1 (in increments of 0.05), and of low severity patients between 0 and 1 (in increments of 0.1).

We found that the driving factor in increasing efficiency was the fraction of high-severity patients seen. This effect is illustrated in Fig. 8, a contour plot of total time in system for all patients vs. the percent of each type of patients seen by residents. The contour lines are all nearly vertical, which shows that the driving factor is percent of high severity patients seen. The reasons are threefold: the majority (75%) of patients in the UMMC ED are high severity, residents have a much bigger effect on the service time for high-severity patients than for low-severity ones (5.3% vs. 1.9%), and high-severity patients take about twice as long to treat (8 h vs. 4 h), so a similar percent reduction in their service time will more heavily influence the average total time in system. We hypothesize that residents increase treatment speed for high-severity patients more than low-severity patients because more complex care is required and there are more chances for work to be done in parallel with attending physicians on high-severity patients. On the other hand, with lower severity patients, the complexity of treatment is lower, and there are fewer chances for work to be done in parallel, so the treatment times are not reduced as much.

The effect of residents on throughput is similar. The percent of high severity patients seen by a resident has a strong effect on throughput, while the percent of low severity patients seen has no detectible effect on throughput. Fig. 9 shows a contour plot of throughput vs. resident presence. Raising the percent of high severity patients treated by a resident from 50 to 100 increases throughput by 2.7%. The contour lines are essentially vertical, meaning that changing the percent of low severity patients treated by a resident has no effect on throughput. Again, this may be because high severity patients take longer to treat, are a higher fraction of the ED patient population, and because residents have a larger effect on their service times.
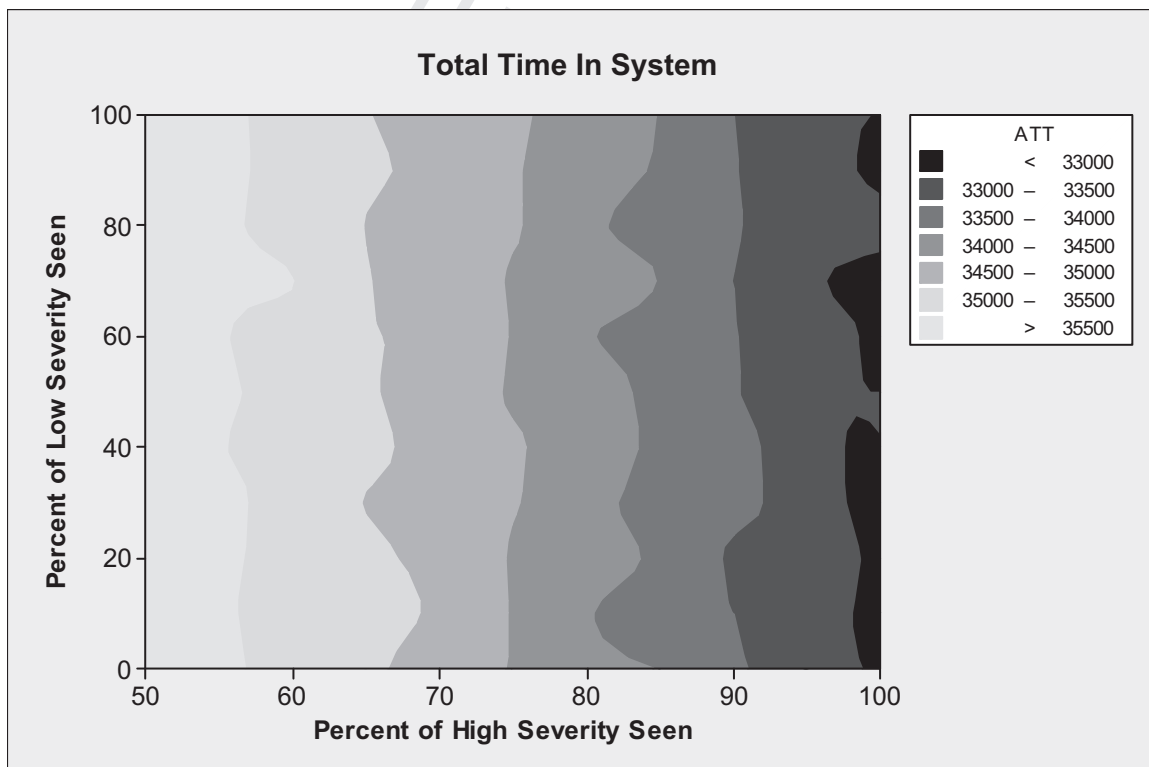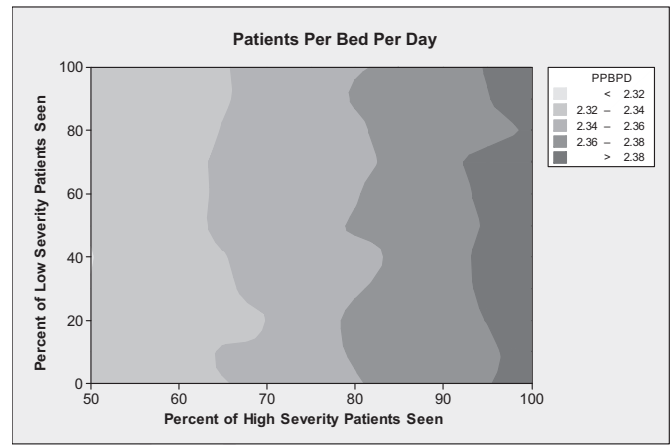

Fig. 9. Contour plot of throughput (in patients per bed per day) vs. resident presence.

In our third experiment, we tested the effect of resident presence on efficiency when treating a variety of patient populations, to see the effects of residents in medical centers that are similar to UMMC but that have different patient characteristics. We generated two additional patient populations, one with a predominantly high-severity patient population (90% high severity), and one with a predominantly low-severity population (50% high severity). All other treatment and patient attributes were held the same. We then looked at the effect of having residents present on efficiency.

We saw that, regardless of patient mix, residents still have a positive effect on system efficiency. Figs. 10 and 11 show the effect residents have on total time in the ED and waiting time, respectively. In both patient populations, residents have positive effects on efficiency. The total time in system effect is about the same for
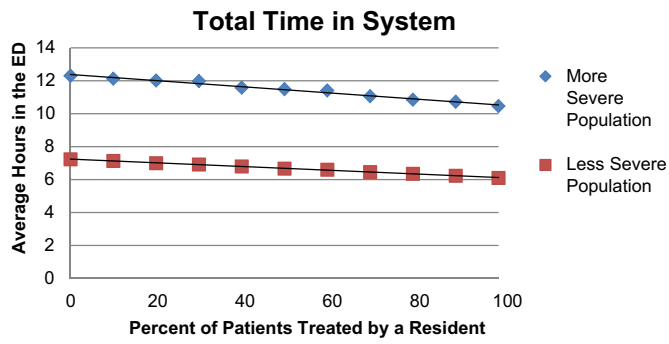

Fig. 8. Total time in system (in seconds) vs. resident presence.

Total Time in System

Fig. 10. Total time in system (in hours) vs. resident presence for different patient severity mixes.

both populations, reducing average time from 12.3 h to 10.5 h in the more severe population, and from 7.2 h to 6.1 h in the less severe population. The lines for the two populations are essentially parallel, meaning that the effect is the same in both patient populations. In both populations, total time in system is reduced by about 15%.

Residents also had an effect on waiting time in both populations. Fig. 9 shows a graph of time to first bed vs. resident presence for the more and less severe populations. In this case, residents had a much more significant effect on first bed time in the more severe population than in the less severe population. This may be because more severe patients take longer to treat, and, therefore, they increase the stress on the system. This leads to longer queues and more variation in waiting time. This means that a similar reduction in processing time has a greater impact on waiting times in the high severity population than in the low severity population.

6. A related queueing model

In addition to the simulation, we also used queueing theory to model the flow of patients through the ED. Specifically, we chose to use an $M/G/k$ queue to represent the system, with each bed being treated as a server. This requires a few simplifying assumptions. First, we assume that patients arrive according to a Poisson process with a fixed arrival rate. Second, we assume that all patients who enter the queue will wait until they are served (no abandonment). Third, we assume that patients are treated in the order in which they arrive (first come, first served). We analyzed the queue with two different service time distributions. The first was the empirical distribution for patients treated on Wednesday mornings, when residents were absent. The second distribution was the empirical
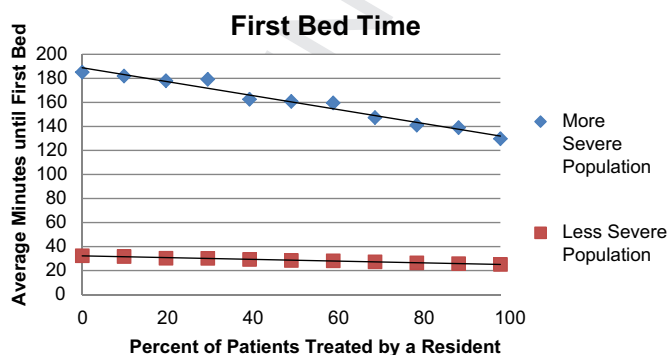
First Bed Time

Fig. 11. Time to first bed (in minutes) vs. resident presence for different patient severity mixes.

distribution for all patients treated during the mornings of the other weekdays (when residents were present), excluding weekends. These simplifying assumptions make the model tractable and allow us to estimate the average waiting times and queue lengths for the system with and without residents.

No closed-form solution to the $M/G/k$ queue exists, so we use the approximations derived by [19]; which take into account the first two moments of the service time distributions. An average of 2.94 patients arrived at the ED per hour and there are 27 beds in the ED. The mean treatment time of patients when residents were absent was 8.22 h and the mean squared treatment time was 121.65 h. When residents were present, the mean treatment time was 7.9 h and the mean squared treatment time was 103.6 h. The queueing model reports that the average waiting time of patients when residents are present is 55 min, compared to 135 min when residents are absent. So, when residents are present, we observe a 59% reduction in waiting time. The residents have a similar effect of the average number of patients in the waiting room (average queue length). If residents were always present, the average queue length would be 2.7, compared to 6.6 if residents were always absent, again a reduction of about 59%.

The waiting times predicted by the queueing model are lower than the historical averages, as a result of the simplifying assumptions. The queueing model has less variability than the real system, so it will have fewer occurrences of high congestion, which leads to lower average waiting times. While the waiting times are smaller, the effect that residents have on waiting times is a 59% reduction in the queueing model compared to 35% in the simulation. Although the queueing model cannot address all of the questions that the simulation can, with respect to time to first bed, the two models at least point in the same direction. This serves to enhance our confidence in the simulation model.

7. Conclusion

A common hypothesis in the medical community is that residents slow down treatment in EDs and have a negative impact on system efficiency, compared to just attending physicians. This paper has shown that, to the contrary, residents have a positive effect on throughput and treatment times. In particular, we found that, when treating high severity patients, residents help to decrease waiting times, decrease treatment times, and increase throughput. While efficiency might not be a main concern in deciding which patients are seen by residents, we would recommend that they see as many high severity patients as is feasible. This fits with the mission of the ED residency program. Furthermore, since residents cannot work as many hours per week as in the past, it is important for them to use their time wisely and productively. The main contribution of this paper is to provide evidence refuting the hypothesis that residents slow down progress in the ED and that they have a negative effect on efficiency.

References

[1] Brailsford SC, Harper PR, Patel B, Pitt M. An analysis of the academic literature on simulation and modeling in health care. Journal of Simulation 2009;3:130–40.
[2] Ceglowski R, Churilov L, Wasserthiel J. Combining data mining and discrete event simulation for a value-added view of a hospital emergency department. Journal of the Operational Research Society 2007;58:246–54.

[3] Dexter F, Macario A, O'Neill L. Scheduling surgical cases into overflow block time – computer simulation of scheduling strategies on operating room labor costs. Anesthesia & Analgesia. 2000;90(4):980–8.

[4] Dowd MD, Tarantino C, Barnett TM, Fitzmaurice L, Knapp JF. Resident efficiency in a pediatric emergency department. Academic Emergency Medicine 2005;12(12):1240–4.

[5] Eappen S, Flanagan H, Bhattacharyya N. Introduction of anesthesia resident trainees to the operating room does not lead to changes in anesthesia-controlled times for efficiency measures. Anesthesiology 2004;101(5):1210–4.

[6] Fone D, Hollinghurst S, Temple M, Round A, Lester N, Weightman A, et al. Systematic review of the use and value of computer simulation modeling in population health and health care delivery. Journal of Public Health Medicine. 2003;25(4):325–35.

[7] Fritze J. Medical expenses have 'very steep rate of growth'. USA Today 2010. 2/4/2010.

[8] Fomundam S, Herrmann J. A survey of queuing theory applications in healthcare. Working Paper. Digital Repository at the Univeristy of Maryland; 2007.

[9] Gorunescu F, McClean I, Millard PH. A queueing model for bed-occupancy management and planning of hospitals. Journal of the Operational Research Society. 2002;53(1):19–24.

[10] Green L. Queueing analysis in healthcare. In: Hall RW, editor. Patient flow: reducing delay in healthcare delivery. Springer; 2006. p. 281–307.

[11] Green LV, Soares J, Giglio JF, Green RA. Using queueing theory to increase the effectiveness of emergency department provider staffing. Academic Emergency Medicine: Official Journal of the Society for Academic Emergency Medicine 2006;13(1):61–8.

[12] Harvey M, Al Shaar M, Cave G, Wallace M, Brydon P. Correlation of physician seniority with increased emergency department efficiency during a resident doctors' strike. Journal of the New Zealand Medical Association 2008; 121(1272):59–68.

[13] Hilbe J. Logistic regression models. CRC Press; 2009.

[14] Jacobson SH, Hall SN, Swisher JR. Discrete-event simulation of health care systems. In: Hall RW, editor. Patient flow: reducing delay in healthcare delivery. Springer; 2006. p. 211–52.

[15] Jun JB, Jasobson SH, Swisher JR. Application of discrete-event simulation in health care clinics: a survey. Journal of the Operational Research Society 1999; 50(2):109–23.

[16] Keehan S, Sisko A, Truffer C, Smith S, Cowan C, Poisal J, et al. Health spending projections through 2017: the baby-boom generation is coming to Medicare. Health Affairs 2008;27:W145–55.

[17] Lammers RL, Roiger M, Rice L, Overton DT, Cucos D. The effect of a new emergency medicine residency program on patient length of stay in a community hospital emergency department. Academic Emergency Medicine 2003;10(7):725–30.

[18] Lehaney B, Hlupic V. Simulation modeling for resource allocation and planning in the health sector. The Journal of the Royal Society for the Promotion of Health 1995;115(6):382–5.

[19] Nozaki S, Ross S. Approximations in finite-capacity multi-server queues with Poisson arrivals. Journal of Applied Probability 1978;15(4):826–34.

[20] Offner PJ, Hawkes A, Madayag R, Seale F, Maines C. General surgery residents improve efficiency but not outcome of trauma care. The Journal of Trauma 2003;55(1):14–7.

[21] Salazar A, Corbella X, Onaga H, et al. Impact of a resident strike on emergency department quality indicators at an urban teaching hospital. Academic Emergency Medicine 2001;8(8):804–8.