Service Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

The Impact of a Discharge Holding Area on the Throughput of a Pediatric Unit

Robert Saltzman, Theresa Roeder, Judith Lambton, Lila Param, Brian Frost, Roxanne Fernandes

informs
http://pubsonline.informs.org/journal/serv/

SERVICE SCIENCE
Vol. 9, No. 2, June 2017, pp. 121–135
ISSN 2164-3962 (print), ISSN 2164-3970 (online)

The Impact of a Discharge Holding Area on the Throughput of a Pediatric Unit

Robert Saltzman,[a] Theresa Roeder,[a] Judith Lambton,[b] Lila Param,[c] Brian Frost,[c] Roxanne Fernandes[c]

[a] Department of Decision Sciences, San Francisco State University, San Francisco, California 94132; [b] School of Nursing, University of San Francisco, San Francisco, California 94117; [c] Children's Hospital of Minnesota, Minneapolis, Minnesota 55404

Contact: saltzman@sfsu.edu (RS); tmroeder@sfsu.edu (TR); lambtonj@usfca.edu (JL); lila.param@childrensmn.org (LP); brian.frost@childrensmn.org (BF); roxanne.fernandes@childrensmn.org (RF)

Received: July 24, 2015
Revised: August 18, 2016
Accepted: November 12, 2016
Published Online: May 1, 2017

https://doi.org/10.1287/serv.2016.0167

Abstract. Hospital patients often move from one bed to another for both medical and nonmedical reasons. In a highly utilized quaternary inpatient pediatric unit we have studied, bed and nursing resources are stressed not only by frequent movement of patients but also by the unit's patient discharge policy. We present a discrete-event simulation model for examining how the unit's efficiency may be improved by a better discharge policy. In particular, we use the base version of the model to investigate the impact of sending various percentages of discharge-ready patients to a discharge holding area where they can safely wait for a few hours until being picked up by their parent or guardian. Doing so frees up inpatient beds, allowing the unit to serve many more pediatric patients per year. In a revised version of the model we quantify the benefits of helping some patients achieve discharge-ready status a few hours earlier than under current operations. In both cases, our cost analysis shows that the unit could realize hundreds of thousands of dollars more per year in net revenue. This argument can be used to help persuade decision makers, who have otherwise been skeptical of the idea of a discharge holding area.

Keywords: patient flow • simulation • pediatrics • discharge policy • cost analysis

1. Introduction and Problem Definition

Various statistics have been offered on the number of children admitted for hospital stays each year. While most of the admissions were a result of newborns/maternal care, 1.1 million patients were admitted for acute care in 2004 (Merrill and Owens 2007a, b). As the care of premature infants, genetic diseases such as cystic fibrosis, and congenital heart disease has improved, these children often require frequent hospital admission. Compounding the issue is the trend for hospitals to close inpatient pediatric units whose percentage of total hospital admissions is calculated at about 5% (Maguire 2007). California alone has lost more than 800 pediatric beds in the past decade as many facilities have converted those beds to adult services (Poindexter et al. 2014), while many smaller hospitals have converted formerly pediatric-only beds to mixed use beds that can be used for either children or adults. Although the total number and duration of pediatric hospitalizations may be declining overall, the sickest children requiring the most sophisticated care must now seek services more often at regional centers (such as the one studied here) and the burden on these facilities is onerous. With already high and increasing bed occupancy rates, these centers are having difficulty accommodating all pediatric patients.

Patients flow through hospitals in a stochastic fashion that depends on attributes such as their diagnosis, admit type, severity, sex, age, and infectious status. As they do so, patients compete for bed and nursing resources, and may require different amounts of nursing resources at different times. The hospital bed controller (typically the hospital nursing supervisor), in consultation with charge nurses on each floor, constantly reassesses the status of inpatients and must balance their needs against those of newly arriving patients as well as with one another. For example, the bed controller may find it appropriate for a newly arriving high-severity patient to "bump" (displace) an inpatient located in a room close to the nursing station (and whose status is less critical) to a room further from the nursing station. Similarly, the medical condition of a formerly low-severity patient in a room relatively far from the nursing station may deteriorate to the point where the bed controller feels the patient should be moved to a higher-priority bed closer to the nursing station. If the unit is heavily occupied, this may require the simultaneous "swapping" of patients in different beds. Another possibility is that a low-severity patient needs to be "transferred" (for medical or other reasons) to an open bed closer to the nursing station. In short, the bed controller is repeatedly confronted with different patient movement scenarios and must make trade-offs in placing patients in beds at various distances (i.e., beds for patients of different priorities) from the nursing station.

Allocating beds to pediatric patients is further complicated by the fact that patients must satisfy a number of compatibility conditions before they can be placed in a semiprivate room whose other bed is occupied. For instance, male and female pediatric patients may only share a room if both are under two years of age. There is also greater concern with pediatric patients about their infectiousness. In fact, of the top 10 reasons for pediatric admission, infections of the lung and airways (number one) and infectious diseases (number 10) ranked among the most common (Clark 2011). The infectious nature of these admissions requires not only locating an appropriate bed, but also possibly rearranging children already admitted to accommodate appropriate isolation strategies. In particular, infectious pediatric patients are permitted to share a room only if both patients in the room are infectious, at least five years old, and of the same sex. These requirements add complexity to patient flow as children may have to be moved across inpatient units. Thus, pediatric patients can only be admitted to the hospital if the bed controller can find a *suitable* bed. If none is available, new patients must wait in some waiting area for a suitable bed to open up. This increases their medical risk as they are not under professional supervision. Those who find the wait too long may leave the hospital without being seen.

The bed controller faces many challenges that stem from a high bed occupancy rate. While hospitals with relatively low bed occupancy normally have enough beds to satisfy newly arriving and current patients, bed controllers in highly utilized hospitals, such as the quaternary (i.e., tertiary, with high levels of specialization) facility we have studied, have a scarcity of open beds to work with and cannot accommodate a substantial portion of new arrivals and inpatient movement requests. Consequently, we wished to investigate whether a strategy that does not require major structural changes to the hospital could be employed to make it more likely that inpatient beds might be available more often.

In particular, the hospital under study had previously considered the use of a discharge holding area (DHA) for patients, but administrators were unable to clearly see the financial benefits of such a policy, so never implemented the concept. Our primary goal was to better understand the bed allocation process and demonstrate how a DHA for pediatric patients might be able to not only free up limited resources and thus reduce resource blocking, inpatient movement, and patient waiting times, but also benefit the hospital financially. We have not seen research addressing this issue.

To do so, we built a discrete event simulation (DES) model, "a good choice when many patient characteristics must be taken into account, particularly if they change over time," and when "patients' demand for particular resources and their priority status in a queue may be influenced by their attributes" (Karnon et al. 2012, p. 703).

The next section reviews relevant literature while Section 3 presents our simulation model in detail. Data sources and key model inputs are the subject of Section 4, and Section 5 describes steps taken to validate the model and performance measures used to evaluate its output. Section 6 presents our experimentation with the model, numerical results, and related discussion. The article concludes in Section 7 with some directions for future work.

2. Relevant Literature

Computer simulation has been used extensively by operations researchers to examine issues in healthcare. Jun et al. (1999) survey 117 articles applying DES to healthcare problems, organizing them into various categories of healthcare, such as patient scheduling and admissions, patient routing and flow schemes, and scheduling and availability of resources. Katsaliaki and Mustafee (2011) survey 201 articles from 1970–2007 that apply simulation to healthcare; the majority use Monte Carlo simulation, mostly to assess health risks from exposure to certain drugs, and for economic modelling of healthcare interventions.

One of the challenges we faced in our study was the availability of patient move data, which appears to be a common problem. One of the more sophisticated approaches we found was Connelly and Bair (2004), who use DES to compare two types of emergency department triaging, fast-track and severity ratio. Data were collected for a five-day period from electronic records and a manual review of 674 emergency department charts. Study results were mixed. While the severity ratio approach improved treatment and service times for higher-severity patients, low severity patients experienced an increase in both times. The model relied on manual chart reviews for almost 600 patients, making the data realistic but very time-consuming to collect.

In a large, multidisciplinary project at the Veterans Administration Pittsburgh Healthcare System, Bountourelis et al. (2011, 2012) develop a complex DES model to simulate patient flow through the intensive care unit (ICU). By analyzing data from bed management and patient movement databases, they were able to accurately estimate and incorporate actual and medical lengths of stay, patient blocking, and transition probabilities into their model. Ceglowski et al. (2007) take an unusual approach to the data gathering phase of their simulation study. Rather than relying on direct observation and interviews, they use a data mining technique that allows the data to self-form into groups of "similar" patients based on treatment. These patient groups are then used in a high-level model that allows a systems-level view of the emergency department, to complement more detailed (traditional) DES models.

While simulation has been used often to study adult patient flow in hospitals, few simulation-based articles appear to address issues specific to pediatric patients. Pediatric patients tend to be more complex to model because of the additional constraints for placement and care. An early paper by Romanin-Jacur and Facchin (1987) describes simulating a pediatric ICU in order to find the number of beds and nurses required to satisfy all admissions and assistance requests. Hung et al. (2007) employ DES to examine the effects of different staffing levels on pediatric emergency department operations, specifically patient flow and waiting times, staff efficiency, and morale. Their model, based on extensive staff interviews and direct observation of over 500 patients, shows that even relatively minor changes in staffing can positively impact patient waits. This matches our results. Ryckman et al. (2009) report on a large-scale project undertaken at a pediatric hospital in Cincinnati whose goal was to match patient demand with available capacity in order to reduce waiting times for urgent care children and avoid the typical long lists of add-on patients at the end of the day. The availability of data was a big challenge to the team, and much of the baseline data for flow through the operating room and pediatric ICU had to be collected by hand.

Although their study does not deal with pediatric patients, Crawford et al. (2014) use DES to examine the impact of discharge timing on the performance of an acute care hospital. They find that two proactive strategies that consider the number of patients waiting for emergency department (ED) beds when making the decision to discharge a patient significantly lower ED waiting and boarding times but increase the number of readmissions. The discharge policy considered here, by contrast, focuses on moving discharge-ready pediatric patients out of inpatient beds to a DHA where they are monitored by nursing staff until being picked up by a parent or guardian. We do not address possible readmission after discharge.

Apart from simulation, queueing theory has also been used to model stochastic networks. While queueing models may be easily applied in some settings, complex situations such as hospital systems can become intractable to represent as queuing models. As described in Cochran and Bharti (2006), queueing networks may be most valuable in the initial modeling stages, to understand in broad strokes the nature of the system, to communicate to stakeholders, and to provide high-level guidance in system design. After the initial stage, the authors use DES to model the system at a detailed level to fine tune the initial bed balancing achieved through the queueing network.

As an alternative to simulation and queueing theory, Thompson et al. (2009) develop a finite-horizon Markov decision process to assist in proactively moving patients between floors in anticipation of surges in demand. Because their model is too large to solve for practical use, the authors develop an approximate decision rule system. The resulting decision support system was implemented at Windham Hospital in Connecticut, with promising initial results.

3. Simulation Model in Detail

Since we wanted to better understand the hospital's bed allocation process, visualize its operations, track a variety of performance measures, and test out potential new discharge policies, developing a detailed, animated computer simulation model seemed the most appropriate approach. Our model, referred to here as the bed assignment and patient transfer simulation model (BAPTSM), was built using Rockwell Automation's Arena v14 simulation software package (Kelton et al. 2010). For concreteness and ease of presentation in the rest of the paper, we focus on activity within the Pediatric Acute-Medical/Surgery (PAMS) unit, the largest of the hospital's eight pediatric units, accounting for 41% of all admitted patients between July 2005 and June 2011; however, it would not be difficult to adapt our analysis to any other unit.

3.1. Description

The simulation model is driven by the arrival of new patients (entities) who are assigned attributes that affect their experience in the hospital, including their primary diagnosis class, admit type, bed priority, sex, age, infectious status, and medical length of stay (LOS). To approximate the mental model employed by the bed controller and charge nurses, who give priority to placing higher-severity patients in beds closer to the nursing station, we divide the unit's beds into three groups: A, B, and C (see Figure 1). The A beds are found in private rooms nearest the nursing station; B and C beds are in semiprivate rooms located further from the nursing station, with B beds being slightly closer to the nursing station than C beds.

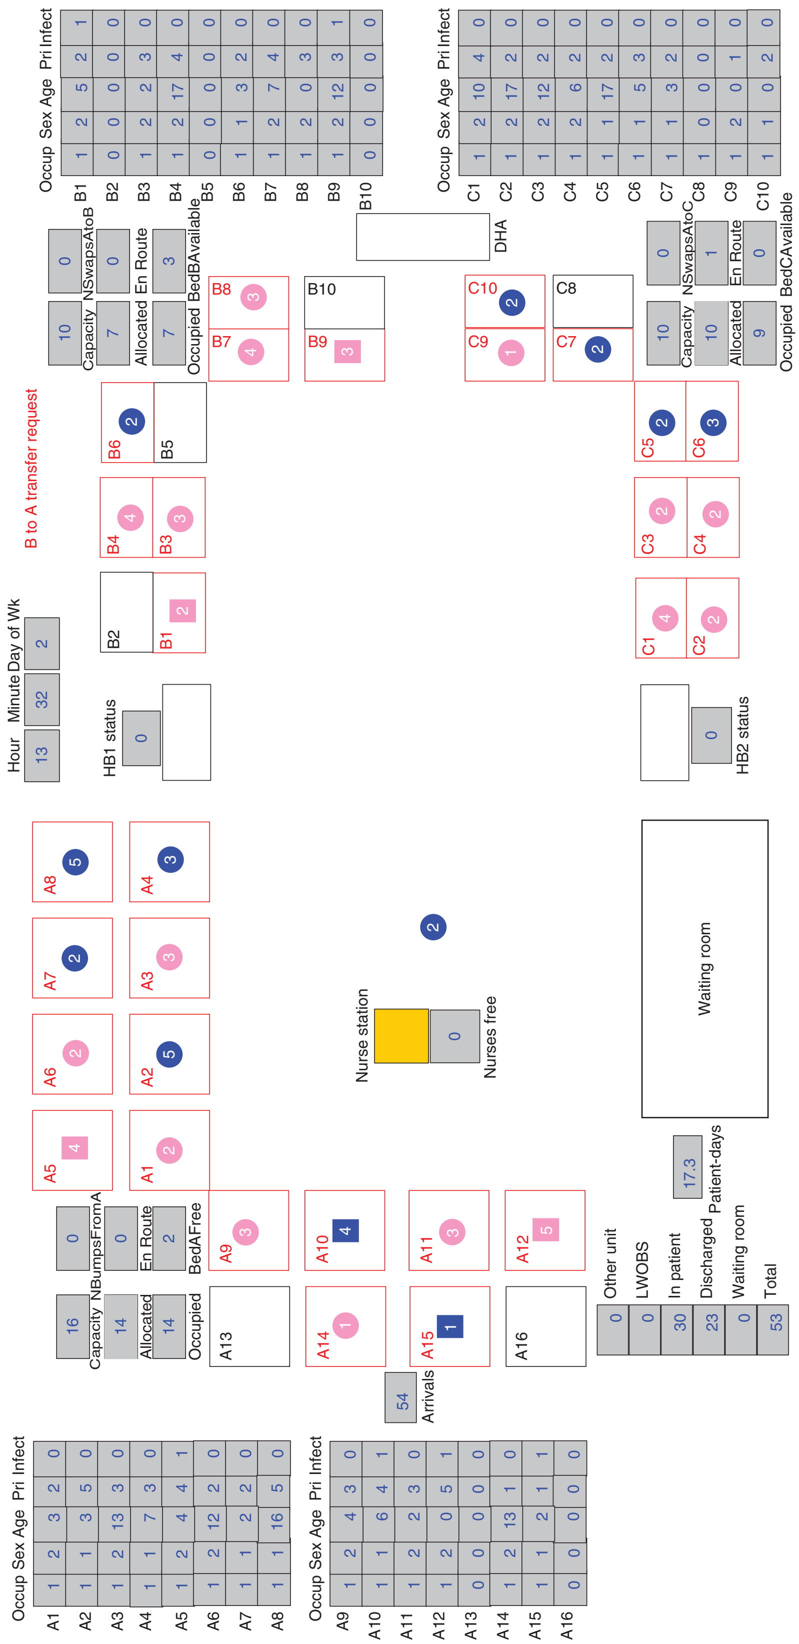

Figure 1. Snapshot of the Animated Portion of BAPTSM

Table 1. Initial Bed Priority in the Model for PAMS Patients, 5 = Highest Priority

Class	National hospital discharge survey: 2007 summary category of first-listed diagnosis	ICD-9-CM	Admit type				PAMS unit (%)
			Emerg	Urgent	Transplt	Routine	
1	Infectious and parasitic diseases	001–139	4	4	3	3	0.8
2	Neoplasms	140–239	3	3	2	2	1.5
3	Endocrine, nutrit'l and metabolic diseases, immunity disorders	240–279	2	2	1	1	5.2
4	Diseases of the blood and blood-forming organs	280–289	3	3	2	2	0.8
5	Mental disorders	290–319	3	3	2	2	0.4
6	Diseases of the nervous system and sense organs	320–389	4	4	3	3	8.7
7	Diseases of the circulatory system	390–459	4	4	3	3	2.0
8	Diseases of the respiratory system	460–519	4	4	3	3	7.7
9	Diseases of the digestive system	520–579	2	2	1	1	10.2
10	Diseases of the genitourinary system	580–629	2	2	1	1	6.6
11	Complications of pregnancy, childbirth and the puerperium	630–677	2	2	1	1	0.0
12	Diseases of the skin and subcutaneous tissue	680–709	2	2	1	1	2.0
13	Diseases of the musculoskeletal system and connective tissue	710–739	2	2	1	1	4.2
14	Congenital anomalies	740–759	3	3	2	2	12.2
15	Certain conditions originating in the perinatal period	760–779	3	3	2	2	0.6
16	Symptoms, signs, and ill-defined conditions	780–799	4	4	3	3	24.1
17	Injury and poisoning	800–999	4	4	3	3	2.1
18	Supplementary classifications	V01–V86	4	4	3	3	10.7
			20.4%	38.3%	1.6%	39.7%	100.0

The initial search for a suitable bed in the model depends on the patient's bed priority, a proxy for the patient's actual severity. Because patient severity levels were not given in the data sets available to us, we estimate a patient's bed priority on a scale of 1 to 5 (with 5 indicating the highest priority for a bed) based on two other attributes for which we did have data, namely, major diagnosis class and admit type. The International Classification of Diseases, Ninth Revision, Clinical Modification (ICD-9-CM) is the official system of assigning codes to diagnoses and procedures in U.S. hospitals (CDC 2015). To simplify, we condensed the thousands of ICD-9-CM codes into 18 major classes, as in Hall et al. (2010).

Admit type, by contrast, is a nonstandard code used by the particular hospital under study to indicate one of four patient types: emergency, urgent, transplant/donor, routine/elective. The percentages of patients falling into each ICD-9-CM class and admit type category for the hospital's PAMS unit are shown in the right and bottom margins, respectively, of Table 1. In the model, these two attributes determine the patient's initial bed priority (see columns 4–7 of Table 1) and are based on the judgment of several coauthors from many years of clinical experience.

BAPTSM logic approximates the initial bed allocation process as follows. For patients with a bed priority of 1 or 2, the bed controller first tries to find a suitable C bed; if none is available, the controller looks for a suitable B bed; otherwise, the controller seeks an available A bed. For patients with a bed priority of 3 or 4, the controller first tries to find a compatible B or C bed; if none is available, the controller looks for an open A bed. Finally, for patients with a bed priority of 5, the controller first seeks an open A bed; if none is available, an existing A-bed patient with lower priority is "bumped" (moved out) to make room for this high-priority patient. If no suitable B or C bed can be found for the bumped patient, the controller sends the patient to another unit within the hospital. While this final move may not always match the real bed allocation policy, we feel that this assumption is an adequate representation of the system for the purposes of this study.

In addition to bumps, BAPTSM handles two other special kinds of inpatient moves faced by the bed controller: transfers and swaps. During transfers, higher priority patients move to a bed nearer the nursing station where they can be monitored more closely. In the model, patients in B (or C) beds may be transferred to A (or B) beds if one becomes available before their medical LOS is completed. Swaps are more complicated situations in which patients in B (or C) beds need to move to an A bed (e.g., because their condition has deteriorated) but none is currently available. To accomplish this, patients in both bed groups are temporarily moved to hall beds, and eventually on to the desired A and B (or C) beds, respectively. The model simulates two-way swaps such as this, but not three-way swaps.

Prior to placing a patient in a B or C bed, the model checks patient compatibility conditions described earlier. Patients for whom a suitable bed is found are admitted to the hospital and then routed from the nursing station to a specific bed and room. If none is found, patients are instead sent to the waiting room where they wait for a bed. However, patients not admitted to the hospital within a few hours abandon the waiting room and leave the hospital without being treated.

Besides beds, BAPTSM also accounts for nursing resources. While patient arrivals fluctuate hourly, the size of the nursing staff remains essentially constant throughout the day, with the same number of nurses working each of the two 12.5-hour shifts. This is because nurses must constantly monitor existing patients, even as new ones are admitted. In a pediatric unit such as this one, nursing standards permit each nurse to care for only three or four patients, depending on patients' age and severity. Since this hospital's PAMS unit has 36 beds, the model assumes a staff of 10 nurses.

The third author conducted a survey of hospital staff members and found that the three special kinds of inpatient moves require nursing staff and take anywhere from 30 to 90 minutes to complete. Furthermore, each inpatient move requires at least two nurses for a minimum of 20 minutes, plus one nurse for the remainder of the move. In practice, additional resources such as a unit clerk and the bed control person may also be involved in these moves for brief amounts of time, but the model does not track these resources.

More routine patient movements in BAPTSM are assumed to take a shorter amount of time (e.g., 10 minutes) and require less staff than the special kinds of movement. For example, patients admitted upon arrival require only one nurse to move from the nurse station to their room, while those not admitted walk unaided to the waiting room.

As in practice, a vacated bed in the model must be cleaned before another patient may occupy it. Based on her experience, the third author estimated that bed cleaning takes approximately 25 minutes. During this time, new patients may not occupy the bed; we assumed that resources required to do the bed cleaning do not constrain hospital operations and, therefore, BAPTSM does not track them.

Once pediatric patients finish their medical LOS, they are "discharge possible" (see Figure 2), but not actually discharged from the hospital. Pediatric patients may have to spend more than another full day in the hospital after becoming discharge possible. In general, they require more time to be discharged than adult patients because, for instance, the admitting physician must examine the child before she can become "discharge ready," and physicians only see their patients in this hospital during morning (7:00–9:00 A.M.) and evening (5:00–7:00 P.M.) rounds. Discharge-possible patients seen by their physician during morning rounds become discharge ready later that same day after receiving a host of ancillary services such as pharmacy medications, rehabilitation requests, social service, and nutritional consults; those seen during evening rounds become discharge ready the *following* day since the required ancillary services are unavailable after 5:00 P.M. Furthermore, even when discharge-ready, pediatric patients cannot be released from the hospital on their own; as minors, they must be released to a parent or guardian who may be away from the hospital at this time. Consequently, pediatric patients often spend a good deal more time occupying beds than is medically necessary.

To simulate this final phase of the hospital stay, patients in the model receive values for three more attributes: (1) a discharge-possible time (*DPTime*), the time at which the patient's randomly generated medical LOS ends; (2) a discharge-ready time (*DRTime*), a randomly generated time between 2:30 P.M. and 3:30 P.M. either the same day as the day on which *DPTime* occurs (if medically cleared by the physician during *morning* rounds) or the next day (if medically cleared during *evening* rounds); and (3) a "discharge-actual" time (*DATime*), the time at which a parent or guardian actually picks up their discharge-ready patient from the hospital. *DATime* is assumed to take on a random value between 5:00 P.M. and 9:00 P.M. (i.e., after the parent gets off work and reaches the hospital). While some parents may be able to pick up their patients immediately after they become discharge ready, the majority of parents of patients at this public, urban, quaternary hospital do not have that much flexibility.

Once discharge possible and medically cleared by their physician, patients experience a discharge-ready delay until their *DRTime*. Discharge-ready patients then undergo a final discharge delay, waiting in their rooms until being picked up later that day at *DATime.* In essence, the hospital currently discharges pediatric patients to the

Figure 2. Timeline Illustrating Key Pediatric Patient Events Near the End of Hospitalization

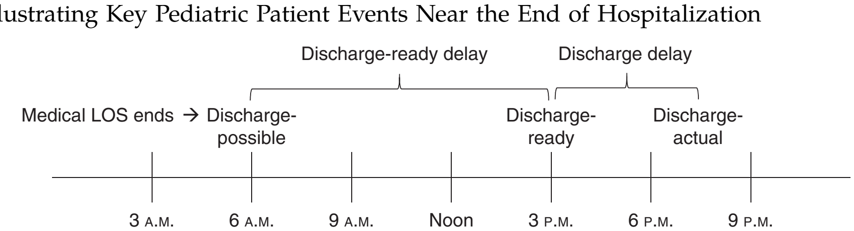

Note. The discharge-possible time can occur at any time of day.

beds they already occupy. By contrast, the model permits discharge-ready patients to be discharged to a DHA, where they wait to be picked up. In the meantime, they no longer occupy an inpatient bed. Freed-up beds can be used to accommodate new patients or help current inpatients get the beds they need.

3.2. Additional Model Assumptions

While trying to accurately portray the behavior of an actual pediatric unit, we make the following additional assumptions to simplify the many complex features of a real hospital.

1. The higher a patient's priority level, the more closely the patient should be monitored. Therefore, each time a patient is "promoted" to a more desirable bed in the model (e.g., from a C to a B bed), the patient's bed priority level increases by 1; conversely, each time a patient is "demoted" to a less desirable bed (e.g., from an A to a B bed), bed priority decreases by 1.

2. An inpatient attempts to move to a new bed whenever the time that has accumulated among inpatients since the last move attempt (measured in patient days) exceeds the model parameter TBMR (time between move requests). The type of moved attempted (bump, swap, or transfer) is random. We select a patient in a given group of beds for promotion (demotion) to a higher (lower) ranked bed based on the highest (lowest) bed priority level, consistent with the notion that patient condition drives most actual bed allocation decisions in the hospital. We break ties among patients with the same priority by selecting the patient occupying the lowest-numbered bed.

3. Inpatients who have been moved from their beds (because of more urgent demands for them) but for whom another suitable bed within the unit cannot be found are sent to another unit in the hospital and no longer tracked by the model.

4. The waiting room has unlimited capacity in the model. In reality, patients may wait in several different locations, e.g., the emergency room, various clinics, and doctors' offices. BAPTSM includes a single waiting room to simplify its layout and flow.

5. Patients in the waiting room who are not admitted within x hours of arrival leave without being seen (LWOBS), where x is uniformly distributed between two and six hours.

6. In reality, a small percentage of patients who are less than one-year-old acquire infections after being admitted to the hospital and require immediate isolation; however, BAPTSM does not capture this relatively infrequent situation.

7. Because of a lack of data regarding the detailed activities of the nursing staff, the model simply keeps nurses occupied 90% of the time with randomly generated generic tasks of random duration; during the remaining 10% of the time, nurses are available to move patients. In our Arena model, patients needing to be moved are held at Seize modules requesting two nurse resources, which requires both nurses to be free before allocating either one to the patient to conduct the move. Thus, patients must often wait to be moved, even when a bed is available, because there are not enough nurses to conduct the move.

4. Data Sources and Key Model Inputs

Numeric data for this project come primarily from two data files from a quaternary pediatric hospital in northern California (with patient identifying information removed or obscured), as well as the 2012 Kids' Inpatient Database (KID), which is based on a nationwide sample of patients at U.S. pediatric hospitals and available from the Agency for Healthcare Research and Quality (AHRQ 2015).

The first hospital file contains 26,166 records of patients admitted between July 2005 and June 2011. Among its fields are the patient's age, sex, admit date and time; admitting department, service, and source; admit type (e.g., emergency, routine); admit diagnosis code (ICD-9-CM) and description. However, it contains no information on length of stay, final diagnosis, or patient arrival (as opposed to admit) time. Using filtering and tabulate features of the JMP v11 statistical software (2013), we were able to determine parameters and distributions for several key simulation model inputs: arrival rates by hour and several patient attributes (sex, age, major diagnosis class, and admit type).

Figure 3 shows the arrival pattern of patients of each admit type, as well as the overall arrival rate; the percentages of patients falling into each admit type every hour fluctuate throughout the day. Because most patients are "unscheduled," we assumed that patients arrive to the unit at random, as in Paul and MacDonald (2013), with a mean that changes hourly. The distribution of admit types varies by hour; we used discrete probability distributions to represent this in the model and assign each patient an admit type.

The sex distribution is stationary, with 52.5% and 47.5% of patients being male and female, respectively. The age distribution differs somewhat for males and females, especially for patients less than one-year-old, so the model uses a different discrete distribution to assign a patient's age based on their sex. We assign a patient's

Figure 3. Admits per Hour to PAMS Unit, by Admit Type, from July 2005 to June 2011

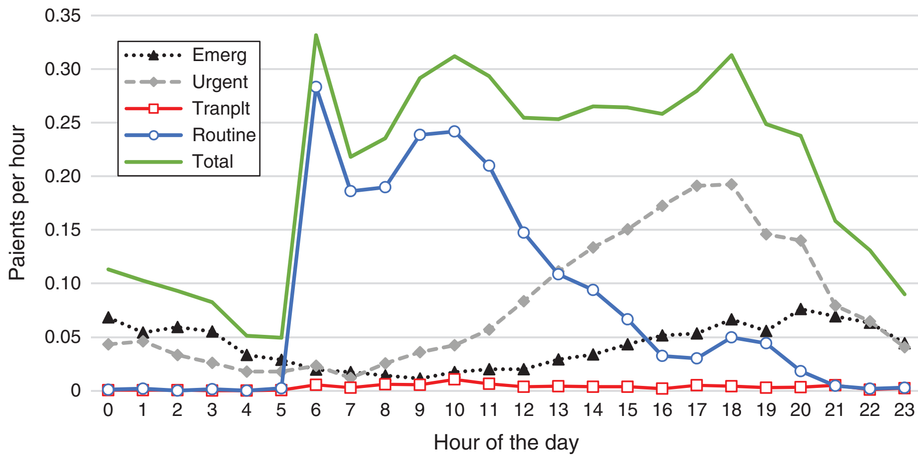

major diagnosis class based on the discrete distribution shown in the right-most column of Table 1, independent of the patient's age or sex. As described earlier, we use the patient's admit type and diagnosis class to assign an initial bed priority, where 1 corresponds to the least urgent and 5 the most urgent cases.

Working with the 2012 KID, we analyzed length of stay data only for patients in children's hospitals (12.5% of its 3.2 million records) since the hospital in our study is also a children's hospital. By filtering and tabulating we found the total LOS mean and standard deviation (in days) of each of the 18 major diagnosis classes (third and fourth columns of Table 2). We then examined the raw data one class at time using Arena's Input Analyzer (Kelton et al. 2010) to identify good-fitting theoretical probability distributions from which BAPTSM samples to assign patients their medical LOS. (We use total LOS as a stand-in for medical LOS since only the former is reported in the 2012 KID.) Because the actual distributions are all strongly skewed to the right and the sample sizes are large, the fitted distributions are not good statistically, but all fit the data well visually. Each theoretical distribution used in BAPTSM is truncated at the maximum observed value for its class (fifth column of Table 2).

The second hospital file contains daily aggregate census data for patients in the hospital between July 2008 and June 2009. It also records each "service day" that a patient spends in a given unit during the same time period, including admit and discharge dates, and room numbers. From these records we could determine patient lengths of stay (but not link them to specific patients in the first hospital file). Overall, we found that PAMS patients had an average total LOS of 6.36 days (or 152.62 hours), a value used to validate model output in the next section.

Table 2. LOS Distributions by Major Diagnosis Class for U.S. Children's Hospitals in 2012

Major class	Actual length of stay				Theoretical probability distribution used in BAPTSM
	n	$\bar{x}$	s	max	
1	15,415	5.37	10.07	265	−0.001 + LOGN(7.61, 19.9)
2	9,569	11.30	17.67	351	−0.001 + LOGN(14.5, 35.9)
3	19,450	4.59	9.00	303	−0.001 + EXPO(4.59)
4	14,784	4.87	8.93	210	−0.001 + LOGN(6.4, 13.2)
5	11,600	7.17	11.02	344	−0.001 + LOGN(10.4, 21)
6	25,007	4.50	9.82	335	−0.001 + EXPO(4.5)
7	7,353	7.95	17.26	358	−0.001 + LOGN(10.9, 31.3)
8	76,310	4.07	8.42	350	−0.001 + EXPO(4.07)
9	36,024	4.79	9.00	365	−0.001 + EXPO(4.79)
10	13,982	3.55	5.63	238	−0.001 + EXPO(3.55)
11	900	3.16	4.31	73	−0.500 + LOGN(3.51, 1.92)
12	12,309	2.98	3.77	162	−0.001 + WEIB(2.36, 0.982)
13	13,519	4.43	6.46	220	−0.001 + EXPO(4.43)
14	30,359	10.42	23.08	356	−0.001 + LOGN(11, 27.8)
15	28,870	16.79	29.50	362	−0.001 + LOGN(23.2, 86.3)
16	23,662	3.39	5.83	252	−0.001 + EXPO(3.39)
17	37,556	4.48	9.65	358	−0.001 + EXPO(4.48)
18	24,125	6.80	13.44	343	−0.001 + LOGN(8.33, 19.4)

By examining the rooms occupied by each patient we could also infer the number of times patients moved within each unit; the average number of intraunit moves per patient (0.255 for PAMS) was used to validate the model. Unfortunately, no data on patient movement *between* units are given in this file.

The file also lists the number of beds in private rooms for each unit (16 out of 36 for PAMS), an important model input; and the average bed occupancy by unit, which is the average daily census of patients (taken at midnight) divided by the total number of beds in the unit. Average bed occupancy for PAMS was 81.9%; this value was also used to validate the model.

An additional input parameter, the probability of an admitted pediatric patient being infectious, was estimated at 30% from Merrill and Owens (2007b). We adjusted a few other input parameters during the model tuning process, as described in the next section.

5. Model Validation and Performance Measures

In addition to using valid data, we took other steps to validate our simulation model, including code validity (verification), face validity, and operational validity.

5.1. Code Validity

During model execution, we animate many activities on screen, helping us verify that the simulation behaves as intended. Figure 1 provides a snapshot of the model. We depict patients with various attributes with different colors, shapes, and numbers as they move around, e.g., male patients are blue, females pink; infectious patients are squares, noninfectious patients are circles, etc. One can also see patients waiting in various places, such as the waiting room and nurse station. We represent each bed resource individually as a black rectangle; when occupied, its outline changes to red; while being cleaned, its shape is filled blue. Numerical displays on the layout show key attributes of each patient occupying a bed, making it possible to confirm that we correctly implemented bed allocation rules.

5.2. Face Validity

The simulation model was developed by the first two authors over a period of several months in consultation with the third author, who has deep system expertise but little experience with operations research models. During this time, the third author shared her knowledge of how pediatric units operate and how she made bed allocation decisions while serving as bed controller at the facility under study. As the model grew in sophistication, the third author repeatedly examined the animated portion and indicated that it realistically portrayed pediatric patient flow and bed control.

5.3. Operational Validity

BAPTSM tracks several performance measures to evaluate the behavior of the hospital unit. Table 3 lists those outputs used primarily to validate the simulation model, those in Table 4 to compare different discharge policies. All are outputs from the model, not inputs to it. For example, the number of admitted patients depends not only on the time-varying input arrival rate, but also on the unit's bed occupancy, which itself depends on the dynamic interactions among many factors including patients and their attributes, nurse availability, and move requests. Not all arriving patients are admitted since the unit's high utilization causes some to wait so long that they leave without being seen. Total LOS depends on not only the input medical LOS distributions, but also the waiting that occurs prior to admission, movement within the unit, discharge-ready delay, and discharge delay. Patient move requests are based on the number of patient days that have accumulated since the last move request; however, the rate at which patient days accumulate depends on the dynamics of the unit. Furthermore, not all move requests are fulfilled since they require both a suitable bed and two free nurses to perform the move. If both conditions cannot be met simultaneously before the patient's medical length of stay ends, the move request is not fulfilled.

Table 3. Performance Measures Used Primarily to Validate the Model

Admitted/Year	Number of patients admitted to the hospital per year
Total LOS	Total time patients actually spend in the unit; includes medical LOS, waiting and transfer time, discharge-ready and discharge delays.
Bed occupancy	Percentage of beds occupied by patients at midnight census.
Moves per patient	Number of times an inpatient is moved within the unit.

Table 4. Additional Performance Measures Used to Evaluate New Policies

Discharged/Year	Number of patients discharged from the hospital per year
Discharge delay	Time between when patients become discharge ready and when they actually leave the hospital, possibly from a DHA.
LWOBS	Number of patients not seen because no suitable bed becomes available during their wait in the waiting room.
Sent to other unit	Number of admitted patients displaced by higher-priority patients.

Table 5. Performance Measures of the PAMS Unit Used to Validate the Simulation Model

Performance measure	Actual PAMS average	BAPTSM base case		Average % difference (Actual – Model) (%)
		Average	95% CI hw	
Admitted/Year (patients)	1,800.0	1,792.2	8.3	0.4
Total LOS (hours)	152.62	152.23	1.10	0.3
Bed occupancy (%)	81.9	82.0	0.2	−0.1
Moves per patient	0.255	0.255	0.002	0.0

Before running the experiments described in the next section, we tuned BAPTSM so that its base case behavior would resemble that of the real unit on the performance measures for which we had actual data. In doing so, we mainly adjusted two input parameters. First, we increased the model's hourly patient arrival rates by 32% to account for the fact that not all patients who arrive actually are admitted to the unit. Second, we set the model parameter TBMR (see Assumption 2) to 22.6 patient days so that the model would generate close to the actual average number of moves per patient.

All simulation model results presented are based on the averages of 100 one-year replications, each of which includes a 90-day warm-up period (excluded from output statistics). Table 5 compares actual PAMS data to the BAPTSM output for the base case (the status quo). For all four performance measures, the model's base case and the actual unit's means are statistically the same at $\alpha = 0.05$, and the percentage differences between them are quite small. Consequently, BAPTSM appears to provide a credible representation of reality.

The average total LOS for the PAMS unit, at 6.36 days, is the second shortest among the eight pediatric units within this hospital. Nonetheless, this unit should be viewed as quite busy, with a bed occupancy rate of nearly 82%, well above the U.S. national hospital average of 66%–67% (Litvak and Pronovost 2010). The unit experiences significant congestion and patient waiting, at times, and is considerably more complex than an ordinary multiserver queueing system operating at 82% utilization. In particular, patient arrivals exhibit substantial seasonality by time of day, with admitted patients varying by type, sex, age, major diagnosis class, and length of stay distribution. Patient moves and bed cleaning, which take time to perform, further reduce the average bed occupancy rate. Finally, because pediatric patients must satisfy a number of compatibility constraints before sharing a semiprivate room, overall bed occupancy is considerably lower than it could be otherwise. Model runs with all beds in private rooms indicate that these compatibility constraints suppress bed occupancy by five–six percentage points.

6. Experimentation

We first experimented with the model to see how sending various percentages of discharge-ready patients to a DHA might impact the unit's performance. While the concept of a DHA is not new (see, e.g., Ashley 1989), it does not seem to have been studied in the context of running a pediatric hospital. One advantage of a DHA is that it need not be located in close proximity to any particular hospital unit; it could be put in any free area and even shared by multiple hospital units. In a pediatric DHA, a nurse and a nurse's aide would be required to monitor all the discharge-ready patients while they wait to be picked up. However, the standard patient-to-nurse ratios would not apply to the DHA since patients located there have already been discharged, that is, medically cleared to leave the hospital.

6.1. Base Model Results

Table 6 shows base model simulated results (means ±95% confidence interval half-widths) for various percentages of discharge-ready patients being sent to the DHA. Performance measures reported in Table 6 include the

Table 6. Base Model Results for Various Percentages of Patients Sent to the DHA

Pct. to DHA	Discharge delay (hours)	Total LOS (hours)	Discharged (patients/yr)	Discharged % Δ	LWOBS (patients/yr)	Sent to other unit (patients/yr)	Bed occupancy (%)
0	4.00 ± 0.01	152.23 ± 1.10	1,729.2 ± 9.2		590.1 ± 11.9	62.9 ± 1.7	82.0 ± 0.2
25	**2.99 ± 0.01**	151.96 ± 1.03	1,740.6 ± 9.9	0.7	577.5 ± 12.1	64.4 ± 1.9	**83.1 ± 0.2**
50	**1.99 ± 0.01**	150.50 ± 1.15	**1,757.0 ± 10.6**	**1.6**	**563.0 ± 12.9**	62.6 ± 1.8	**83.5 ± 0.3**
75	**0.98 ± 0.01**	150.17 ± 1.03	**1,756.7 ± 10.1**	**1.6**	**560.1 ± 12.2**	66.1 ± 1.7	**83.9 ± 0.2**
100	**0.00 ± 0.00**	**148.01 ± 1.08**	**1,771.0 ± 10.2**	**2.4**	**547.8 ± 11.8**	64.0 ± 1.9	**83.9 ± 0.2**

Note. Bolded values are significantly different from the 0% base case at $\alpha = 0.05$.

discharge delay and the total LOS in hours, the annual number of patients who are discharged, LWOBS, and sent to another unit prior to completing their medical LOS, and the bed occupancy rate.

Table 6 reveals that the greater the percentage of patients sent to the DHA, the more efficiently the unit operates: decreases in average discharge delay and total LOS cause the number of discharged (LWOBS) patients per year to rise (fall). Bed occupancy also increases a small but significant amount. Consequently, the unit serves a greater proportion of arriving patients and generates more revenue for the hospital. Although the number of discharges per year in the 25% case does not differ statistically from the status-quo base case (0% sent to DHA), the number of discharges per year at 50%, 75%, and 100% are all statistically higher. While the increases are modest, raising the number of discharges per year by up to 42 patients (2.4%) represents an improvement of practical import (as discussed further in Section 6.2).

Figure 4 shows the average number of patients discharged per year versus the average discharge delay with 95% confidence interval bands at each discharge delay level. The dashed regression line's slope of −9.95 suggests that every hour reduction in the average discharge delay enables nearly 10 additional patients to be treated annually by the unit.

6.2. Revised Model: Reducing the Discharge-Ready Delay

The base case version of BAPTSM reflects the fact that this hospital, similar to others, provides ancillary services (pharmacy medications, rehabilitation requests, social service, and nutritional consults) needed by medically cleared pediatric patients primarily during daytime hours. Consequently, patients cleared for discharge by their physician during *evening* rounds often make little progress toward their actual discharge until the next day. In this experiment, we use the model to test a small process change whereby some ancillary services, i.e., pharmacy medications, can be fulfilled in the evening, rather than only during the day. In particular, we assume that patients medically cleared during evening rounds experience a two-hour drop in their discharge delay so that they are discharge ready between 12:30 and 1:30 P.M. the next day.

Results from the revised model in Table 7 show that taking this fairly simple step could yield additional benefits over the base case. Relative to the status-quo base model (no DHA at all), sending just half of discharge-ready patients to the DHA, along with providing pharmacy services in the evening, increases the number of discharged patients per year by 32 patients (1.9%), and bed occupancy by two percentage points. Sending all discharge-ready patients to the DHA leads to 56 (3.2%) more patients being treated annually by the unit. Financial implications of these results are considered in the next subsection.

Figure 4. Patients Discharged per Year vs. Average Discharge Delay

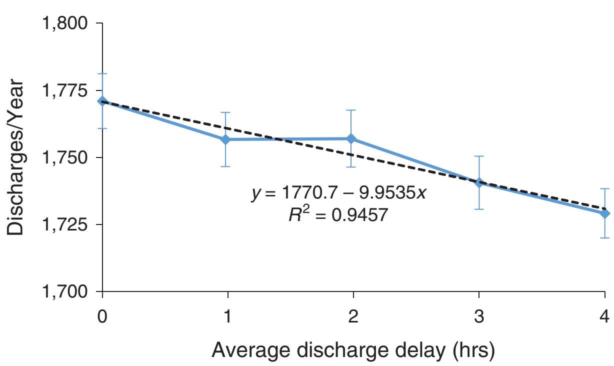

Table 7. Revised Model Results for Various Percentages of Patients Sent to the DHA

Pct. to DHA	Discharge delay (hours)	Total LOS (hours)	Discharged per year (patients)	Discharged per year % Δ	LWOBS per year (patients)	To other unit/Year (patients)	Bed occupancy (%)
Base	4.00 ± 0.01	152.23 ± 1.10	1,729.2 ± 9.2		590.1 ± 11.9	62.9 ± 1.7	82.0 ± 0.2
0	**4.89 ± 0.01**	152.11 ± 1.09	1,734.3 ± 9.8	0.3	585.1 ± 11.7	63.1 ± 1.7	82.1 ± 0.2
25	**3.65 ± 0.01**	151.00 ± 1.08	**1,753.6 ± 10.5**	1.4	566.2 ± 12.4	63.1 ± 1.7	**83.1 ± 0.2**
50	**2.43 ± 0.01**	150.57 ± 1.16	**1,761.9 ± 10.8**	**1.9**	**555.8 ± 13.2**	65.1 ± 1.7	**84.0 ± 0.2**
75	**1.20 ± 0.01**	**149.24 ± 1.16**	**1,776.3 ± 10.8**	**2.7**	**543.0 ± 12. 8**	63.8 ± 1.6	**84.6 ± 0.2**
100	**0.00 ± 0.00**	**147.25 ± 1.06**	**1,785.3 ± 10.8**	**3.2**	**533.4 ± 12.8**	64.7 ± 1.9	**84.7 ± 0.2**

Note. Bolded values are significantly different from the status-quo base case (first row) at $\alpha = 0.05$.

6.3. Cost Analysis

An administrative challenge associated with the DHA is justifying its cost. While inpatients in beds can be billed to insurance companies, time spent in the DHA is not billable, leading to the question of where its funding would come from. In this section, we illustrate that significant additional revenues would likely be realized through the use of a DHA. In doing this analysis, we make the following assumptions in order to conservatively estimate the potential savings.

- The billable daily room rate is $5,189, the actual rate in 2009 (undoubtedly higher now).
- Because both the proportion of private versus public insurance, and the variability in negotiated payment options and rates (case rates, fee for service, percent of charges), are individual to each hospital, we took a conservative approach and assumed the hospital would be reimbursed at 50% of the room charge, excluding all other charges.
- Variable costs (e.g., medication and supplies) are estimated to be 16% (Roberts et al. 1999), so we further reduced the hospital's revenue per patient day by 16% of $5,189. Thus, the hospital's net revenue per patient day is 34% of $5,189, or $1,764.
- The DHA is open daily from (at most) noon–10:00 P.M., or 70 hours per week, and can be staffed by two full-time RNs.
- The annual salary for each registered nurse (RN) staffing the DHA is $75,000.
- Since the hospital is not reimbursed for costs incurred during the patient's discharge delay, revenues are based on the average total LOS minus the average discharge delay (denoted *RLOS*), which varies slightly with the percentage of patients sent to the DHA. We take *RLOS* from the 0% scenario (6.20 days in the base model, 6.17 days in the revised model) so as to compare results across scenarios more consistently.
- The fixed cost of setting up the DHA is not included since we are unable to estimate its value. However, it should be low since little space is required and it need not be medically equipped. It is likely that existing space in the hospital can be repurposed for the DHA.
- The annual cost of keeping the pharmacy open for four hours every evening to fill orders for patients medically cleared for discharge during evening rounds is $85,000.

Table 8. Cost Analysis for Changes in Discharge Policy

Scenario: Model	Scenario: % to DHA	Discharges per year: Ave.	Discharges per year: hw	Annual revenue ($1,000s)	Incremental fixed cost ($1,000s)	Annual net revenue ($1,000s)	Abs. Δ from base ($1,000s)	Rel. Δ from base (%)
	0	1,729.2	9.2	18,911	0	18,911		
Base case	25	1,740.6	9.9	19,036	150	18,886		
	50	**1,757.0**	**10.6**	**19,215**	**150**	**19,065**	**154**	**0.8**
	75	**1,756.7**	**10.1**	**19,212**	**150**	**19,062**	**151**	**0.8**
	100	**1,771.0**	**10.2**	**19,368**	**150**	**19,218**	**307**	**1.6**
	0	1,734.3	9.8	18,967	235	18,732		
Revised	25	**1,753.6**	**10.5**	**19,178**	**235**	**18,943**	**32**	**0.2**
	50	**1,761.9**	**10.8**	**19,269**	**235**	**19,034**	**123**	**0.6**
	75	**1,776.3**	**10.8**	**19,426**	**235**	**19,191**	**280**	**1.5**
	100	**1,785.3**	**10.8**	**19,524**	**235**	**19,289**	**379**	**2.0**

Note. Bolded values are significantly different from the status-quo base case (first row) at $\alpha = 0.05$.

Table 9. Sensitivity of Throughput to Additional Hospital Resources in Base Model

Scenario: Changes in base model	% increase in bed capacity	% increase in nursing capacity	% increase in number of discharged patients
+1 A bed	2.8	0	2.3
+1 A bed, +1 nurse	2.8	10	2.9
+2 A beds	5.6	0	5.1
+2 A beds, +1 nurse	5.6	10	5.3

Table 8 shows the results of the cost analysis, where the italicized scenario in the first row corresponds to the status-quo base case. The average number of patients discharged per year and the 95% confidence interval are given for each scenario. Columns 5, 6, and 7 give the average annual revenue ($Discharges \times RLOS \times$ *Net revenue per patient day*), the annual incremental fixed cost of the two DHA nurses, and the difference between revenue and cost, or net revenue. The last two columns show the difference in revenue between the given scenario and the status-quo base case, when the number of discharges is statistically different, in absolute and percentage terms.

In the base model, half of all discharge-ready patients have to be sent to the DHA to make a statistically significant difference in the number of discharged patients, in which case the unit's annual net revenue is $154,000 (0.8%) higher than the status-quo revenue. Net revenue rises $307,000 per year (1.6%) over the status quo all patients are sent to the DHA when discharge ready.

We experimented with the base model to see if staffing the DHA with two RNs was really the best use of these resources. Model output showed no statistically significant difference in the number of patients treated with the addition of one or even two full-time nurses, demonstrating that beds, not nurses, are the bottleneck resource for this unit. Thus, hiring RNs to staff the DHA appears to have much greater benefit than placing additional RNs in the unit itself.

We also experimented with the base model to see how the unit might be affected by the addition of one or two A beds and possibly more nursing resources (11 nurses available at all times), as seen in Table 9. All scenarios generated statistically significant changes from the base case in the number of discharged patients per year at $\alpha = 5\%$.

Because the unit's demand exceeds its supply of resources, additional beds lead to greater patient throughput, essentially in proportion to the increase in the bed capacity. On the other hand, increasing nursing capacity by 10% helps only slightly, raising throughput by less than a percentage point. One advantage of a DHA is that, unlike a new private room with an A bed, it would require little new infrastructure. Moreover, a DHA can serve multiple hospital units at once, as it will only hold patients who are discharge ready, whereas additional A beds will serve only one unit. Adding two RNs to staff the DHA for 70 hours per week, a 4.2% increase in nursing capacity, will have much greater impact on patient throughput than increasing nursing capacity within the unit itself by 10%.

In the revised model, where the discharge-ready delay drops two hours for patients medically cleared during evening rounds, additional benefits accrue in both the number of discharged patients and net revenue (bottom half of Table 8). Compared to the status-quo operation of the unit, sending 75% of all discharge-ready patients to the DHA allows the unit to treat 47 additional patients annually, generating $280,000 more in net revenue; if all discharge-ready patients are sent to the DHA, the 56 additional patients treated annually increase net revenue by $379,000.

To recap, Table 8 shows that the increase in the annual number of patients discharged by the hospital generates sufficient revenue to easily cover the costs of additional nursing and pharmacy staff to supervise patients in the DHA and keep the pharmacy open for a few hours each evening. Values given here are conservative estimates based on a 50% reimbursement rate; a 75% reimbursement rate would more than double the net revenue increase of each scenario over the base case. In addition, our cost analysis is based on a billable daily room rate from 2009, which is likely to be 30%–40% higher now. Finally, our analysis considers only one the hospital's eight units. If multiple units were to utilize the DHA, additional patient throughput and associated net revenue would be substantial, likely exceeding one million dollars annually.

7. Conclusion

In our review of the literature, we were unable to find papers addressing the flow, and especially the discharge, of pediatric patients. This work was motivated by the desire to serve a larger number of patients without sacrificing quality of care. To operations research practitioners, the proposed solution may seem almost obvious, yet it had been abandoned by the administration of the hospital in question because they were unable to see the financial benefits of the DHA (or, more specifically, how to bill for it).

The results presented in this article make a powerful case for rethinking the way pediatric patients are discharged from the PAMS unit of one particular hospital. We believe similar results would be achieved in other pediatric units. While medical staff may be unwilling to send all discharge-ready patients to the DHA, even sending only half of them can have a significant impact on the number of patients the hospital is able to care for—and thus the revenues the hospital can generate. Although the average time patients spend in the hospital drops by just a few hours (because of briefer discharge delays), the impact over the course of the year is meaningful.

The benefits to patients' well-being come with increased revenues for the hospital. Although insurance companies cannot be billed for time patients spend in the DHA, the hospital's revenues increase because of the increased number of patients served by the hospital over the course of the year. These revenues easily offset the associated personnel costs required to run the DHA. Financial benefits aside, using the DHA would allow this quaternary facility to treat more children who otherwise would have difficulty finding suitable treatment elsewhere.

BAPTSM simulates only the internal functioning of a single unit in the hospital, which necessarily is a great simplification of the overall hospital's operations. Blockages in the PAMS unit can lead to bed allocation problems in other parts of the hospital, as staff are unable to move patients between units. These secondary interactions are not captured in BAPTSM. We assume that patient throughput across the hospital would increase, however, even if only the PAMS unit implements a DHA. Were other units also to adopt its use, greater benefits could be realized.

There do not appear to be major medical risks of using a DHA. Although a small fraction of patients in a DHA will require readmission to the hospital, the same can be said of pediatric patients who have been discharged but still occupy inpatient beds. In both settings, patients have already been medically cleared to leave; readmitting patients from a DHA would be similar to readmitting them from an inpatient bed. In fact, there may be a medical advantage to sending discharge-ready patients to the DHA rather than being kept in their beds because of the increased medical risk of keeping healthy children around patients still being treated.

This project might have benefitted in a number of ways from better and more detailed operational data, such as that contained in an electronic medical record (EMR) system. First, a single EMR database ought to remove some of the inconsistencies we found between our two main data sets, e.g., the differing numbers of admitted patients. Second, EMR data might allow a more accurate determination of patient severity, LOS by severity and service for this hospital, and patient movement between units. Third, EMR data might also provide better insight into the complex, multitasking nature of nursing and how nurses spend their time.

Potential future uses of the model include investigating the impact of varying certain inputs, such as the infectiousness rate and the severity mix of patients, and whether or not managing routine demand to better complement the variations in unscheduled demand would provide the kind of benefits noted by Litvak et al. (2010). Another possibility is to use the model to estimate the benefits of treating all patients in private rooms in pediatric hospitals, an emerging trend in new hospital design (Sadler et al. 2011). BAPTSM could also be employed to quantify how further reductions in the discharge-ready delay could improve efficiency, i.e., what if all ancillary services could be provided in the evening so that pediatric patients who are medically cleared for discharge during evening rounds could actually be discharged that night? Finally, we envision possibly expanding the model to comprise multiple units of the children's hospital, capturing interactions and synergies between closely aligned units, such as the PAMS and Pediatric Acute-Oncology, Bone Marrow Transplant, and Rehabilitation units.

References

AHRQ (2015) Overview of the kids' inpatient database (KID). Accessed May 2, 2015, http://www.hcup-us.ahrq.gov/kidoverview.jsp#about.

Ashley M (1989) Discharge holding area: Using inpatient beds more efficiently. *J. Nursing Admin.* 19(12):32–37.

Bountourelis T, Eckman D, Luangkesorn L, Schaefer A, Nabors SG, Clermont G (2012) Sensitivity analysis of an ICU simulation model. Laroque C, Himmelspach J, Pasupathy R, Rose O, Uhrmacher AM, eds. *Proc. 2012 Winter Simulation Conf.* (IEEE, New York), 931–942.

Bountourelis T, Luangkesorn L, Schaefer A, Maillart L, Nabors SG, Clermont G (2011) Development and validation of a large scale ICU simulation model with blocking. Jain S, Creasey RR, Himmelspach J, White KP, Fu M, eds. *Proc. 2011 Winter Simulation Conf.* (IEEE, New York), 1143–1153.

Ceglowski R, Churilov L, Wasserthiel J (2007) Combining data mining and discrete event simulation for a value-added view of a hospital emergency department. *J Oper. Res. Soc.* 58(2):246–254.

Centers for Disease Control and Prevention (CDC) (2015) International classification of diseases, ninth revision, clinical modification. Accessed August 17, 2016, http://www.cdc.gov/nchs/icd/icd9cm.htm.

Clark C (2011) Top 10 reasons for pediatric hospitalizations. Accessed June 10, 2014, http://www.healthleadersmedia.com/page-1/LED-269728/Top-10-Reasons-For-Pediatric-Hospitalizations.

Cochran JK, Bharti A (2006) Stochastic bed balancing of an obstetrics hospital. *Health Care Management Sci.* 9(1):31–45.

Connelly LG, Bair AE (2004) Discrete event simulation of emergency department activity: A platform for system-level operations research. *Acad. Emergency Med.* 11(11):1177–1185.

Crawford EA, Parikh PJ, Kong N, Thakar CV (2014) Analyzing discharge strategies during acute care: A discrete-event simulation study. *Medical Decision Making* 34(2):231–241.

Hall MJ, DeFrances CJ, Williams SN, Golosinskiy A, Schwartzman A (2010) National hospital discharge survey: 2007 summary. *Natl. Health Stat. Rep.* 29(October):1–20.

Hung GR, Whitehouse SR, O'Neill C, Gray AP, Kissoon N (2007) Computer modeling of patient flow in a pediatric emergency department using discrete event simulation. *Ped. Emergency Care* 23(1):5–10.

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

Karnon J, Stahl J, Brennan A, Caro JJ, Mar J, Möller J (2012) Modeling using discrete event simulation: A report of the ISPOR-SMDM modeling good research practices task force–4. *Medical Decision Making* 32(5):701–711.

Katsaliaki K, Mustafee N (2011) Applications of simulation within the healthcare context. *J. Oper. Res. Soc.* 62(8):1431–1451.

Kelton WD, Sadowski RP, Swets NB (2010) *Simulation with Arena* (McGraw-Hill, New York).

Litvak E, Pronovost PJ (2010) Rethinking rapid response teams. *J. Amer. Medical Assoc.* 304(12):1375–1376.

Litvak E, Gren Vaswani S, Long M, Prenney B (2010) *Managing Variability in Healthcare Delivery* (National Academies Press, Washington, DC).

Maguire P (2007) Getting squeezed: Emerging trends in pediatric inpatient services. Accessed August 17, 2016, http://www.todayshospitalist.com/Getting-squeezed-emerging-trends-in-pediatric-inpatient-services.

Merrill C, Owens PL (2007a) Hospital admissions that began in the emergency department for children and adolescents, 2004: Statistical brief #32. Accessed August 17, 2016, https://www.ncbi.nlm.nih.gov/pubmed/21850769.

Merrill C, Owens PL (2007b) Reasons for being admitted to the hospital through the emergency department for children and adolescents, 2004: Statistical brief #33. Accessed August 17, 2016, https://www.ncbi.nlm.nih.gov/pubmed/21850774.

Paul JA, MacDonald L (2013) A process flow-based framework for nurse demand estimation. *Service Sci.* 5(1):17–28.

Poindexter S, Rañoa R, Parks J, Smith D, Welsh B, Yoshin K (2014) Shrinking services: Pediatric beds in California. Accessed June 10, 2014, http://projects.latimes.com/hospitals/pediatric-beds/.

Roberts RR, Frutos PW, Ciavarella GG, Gussow LM, Mensah EK, Kampe LM, Straus HE, Joseph G, Rydman RJ (1999) Distribution of variable vs. fixed costs of hospital care. *J. Amer. Medical Assoc.* 281(7):644–649.

Romanin-Jacur G, Facchin P (1987) Optimal planning of a pediatric semi-intensive care unit via simulation. *Eur. J. Oper. Res.* 29(2):192–198.

Ryckman FC, Adler E, Anneken AM, Bedinghaus CA, Clayton PJ, Hays KR, Lee B, Morillo-Delerme JW, Schoettker PJ, Yelton PA (2009) Cincinnati Children's Hospital Medical Center: Redesigning perioperative flow using operations management tools to improve access and safety. Litvak E, ed. *Managing Patient Flow in Hospitals: Strategies and Solutions*, 2nd ed. (Joint Commission Resources, Oakbrook Terrace, IL), 97–111.

Sadler BL, Berry LL, Guenther R, Hamilton D, Hessler FA, Merritt C, Parker D (2011) Fable hospital 2.0: The business case for building better health care facilities. *Hastings Center Report* 41(1):13–23.

SAS Institute (2013) *JMP 11: Discovering JMP* (SAS Institute, Cary, NC).

Thompson S, Nunez M, Garfinkel R, Dean MD (2009) Efficient short-term allocation and reallocation of patients to floors of a hospital during demand surges. *Oper. Res.* 57(2):261–273.