PRODUCTION AND OPERATIONS MANAGEMENT

POMS

Vol. 20, No. 3, May–June 2011, pp. ix–xii
ISSN 1059-1478 | EISSN 1937-5956 | 11 | 2003 | 0ix

DOI 10.1111/j.1937-5956.2011.01180.x

Management Insights

An Empirical Study of the Relations between Hospital Volume, Teaching Status and Service Quality
Carol Theokary, Z. Justin Ren

This paper examines the relationships between patient volume, teaching mission, and process quality in U.S. hospitals. Using a large data set that measures process quality for treatments of heart attacks and heart failures in all major U.S. hospitals, the authors find that the impact of patient volume on process quality varies across hospitals with different teaching intensities. Their results suggest that when the complexities of large-scale operations coincide with the added burden of teaching and training functions, teaching hospitals tend to suffer a negative effect on their process quality. The managerial implications of this finding are twofold. First, because large volume hospitals do not always perform better than smaller ones, policy makers and hospital managers should re-evaluate the current practice of regionalization of hospitals services in light of the teaching intensity of the hospital to which services are being redirected. Secondly, while many hospitals define themselves as teaching hospitals and that new doctors must be taught, large teaching hospitals with high resident-to-bed ratios need to address the root causes of the lower quality scores they produced (compared to their smaller counterparts) in order to control quality more tightly.

Performance Effects Related to the Sequence of Integration of Healthcare Technologies
Corey M. Angst, Sarv Devaraj, Carrie C. Queenan, Brad Greenwood

We know that hospitals adopt medical technologies in different sequences. These sequences can result from well planned, long term, system-focused visions, or unstructured market demand-driven processes, or hybrids of the two. If the adopted medical technology proves valuable, it is often integrated into the hospital's information systems. We investigate how the integration sequence impacts hospital performance and found that hospitals which integrated foundational technologies first – which in this case are known to be more complex – exhibit superior performance. While most of the medical technologies in this study have already been adopted by hospitals, many have not yet been integrated into the information system of the hospital and therefore are not interoperable. Because they have not yet been integrated, managers can adjust integration order; our findings suggest that this can impact important performance metrics such as length of stay and costs.

Optimal Design of a Pharmaceutical Price-Volume Agreement Under Asymmetric Information about Expected Market Size
Hui Zhang, Gregory S. Zaric, Tao Huang

Price-volume agreements are commonly negotiated between drug manufacturers and third-party payers for drugs, such as government insurance plans. In one form, a drug manufacturer pays a rebate to the payer on a portion of sales in excess of a specified threshold. We studied the optimal design of a price volume agreement when there is uncertainty regarding the number of units sold, and when both parties (the payer and the manufacturer) have different beliefs about the size of the market. Our results suggest that a payer that is able to negotiate on prices may not always want to include a volume rebate in a contract. In fact, in many cases, the optimal contract includes no rebate term. Common rules of thumb for choosing the rebate value may lead to substantial losses of efficiency.

Access to Long Term Care: The True Cause of Hospital Congestion?
Jonathan Patrick

This research describes a mathematical model for determining the necessary access to long term care (LTC) in order to maintain the number of clients waiting for LTC in the hospital below a pre-specified threshold. The optimal policy was demonstrated to be a threshold policy with priority access given to the hospitals once the census reached a given level. The level triggering priority access was significantly below the prespecified threshold in all instances. Since LTC demand arises in the community as well in the hospitals, the researchers used a simulation model to determine the impact on community client wait times from using the above policy for hospital access. Unsurprisingly, community wait times suffer significantly. While the current situation suggests that no new LTC beds will

be made available, there is a concerted effort through additional home care to allow clients to stay at home longer and therefore hopefully spend less time in LTC. We demonstrated that the average length of stay will need to be reduced by 1/3 in order to achieve the double purpose of maintaining the hospital census below the pre-specified threshold as well as keeping community client wait times below the wait time target of 90 days.

Design and Analysis of Hospital Admission Control for Operational Effectiveness
Jonathan E. Helm, Shervin AhmadBeygi, Mark P. Van Oyen

This paper identifies benefits of an accurate enterprise level management of bed and care resources. US hospitals typically lack this system level management, which contributes to highly fluctuating bed occupancy levels. This system dysfunction manifests itself in emergency patient bed block, cancelation of scheduled surgeries, congestion and operational chaos within the hospital. In congested hospitals, medical patients who are blocked from direct admission to the hospital will tend to use the emergency department as a means to get expedited admission even if it is not a true emergency. Establishing a call-in queue mechanism to serve these patients in an expedited manner can alleviate the strain on the hospital during peak congestion and improve utilization during occupancy dips. By dividing hospital occupancy into three zones (cancelation zone, normal operation zone, and a call-in zone) the hospital can smooth occupancy over time. Implementation involves monitoring occupancy levels, canceling surgeries when the occupancy enters the cancelation zone, and calling in patients from the call-in queue when in the call-in zone. Effective zones can be found using a properly designed simulation of patient flow. A simulation case study with historical hospital data indicates this policy can reduce surgery cancelations and emergency bed block by 15-30%. This admission control mechanism leads to higher quality, better access, and lower cost of care delivery.

Blocking in Healthcare Operations: A New Heuristic and an Application
Kurt M. Bretthauer, H. Sebastian Heese, Hubert Pun, Edwin Coe

Hospital administrators continue to face the challenge of providing the best quality care possible in an environment of limited resources and increasing healthcare costs. Based on data collected from a large hospital system located in the United States, the authors consider the problem of optimal capacity allocation in a hospital setting. Determining the most effective mix of inpatient beds requires taking a system wide view of the hospital and accounting for the complex patient blocking effects between the various units. The authors first present a heuristic for estimating blocking probabilities in an n-stage tandem system that is intuitive, robust, outperforms existing techniques, and is on average within 5% of the exact solution. Next, they develop an extension with general routing that is applicable to a wide range of healthcare and service settings. Applying this method to a real hospital system, the results suggest that changing the number of beds in a downstream unit, such as post-acute care, ripples back through the hospital and impacts blocking probabilities at the upstream units. The findings reinforce the importance of taking into account blocking effects and using a method such as the one presented here to aid in hospital bed mix decisions.

The Surgical Scheduling Problem: Current Research And Future Opportunities
Jerrold H. May, William E. Spangler, David P. Strum, Luis G. Vargas

This paper offers a comprehensive review of the literature to date related to the surgical scheduling process, from the selection of procedures to be performed, to the allocation of resource time to those procedures, and to the sequencing of the procedures within the allocated time. Surgical scheduling is a challenging task, primarily because every detailed-level plan is almost certain to deviate significantly from what actually transpires in the course of the surgical day. Emergency patients enter the system, patients (and perhaps staff) either do not arrive or do not arrive when expected, planned procedures become unnecessary, unplanned procedures become necessary, and procedures take more time or less time than originally planned. In such an environment, a schedule is a guide for operational management rather than a statement of precisely expected outcomes. But the better the guide, the more likely it is that operational management will be able to use resources effectively and efficiently.

Bi-Criteria Scheduling of Surgical Services for an Outpatient Procedure Center
Serhat Gul, Brian T. Denton, John W. Fowler, Todd Huschka

Uncertainty in the duration of surgical procedures can cause long patient wait times, poor utilization of resources, and high overtime costs. The authors provide managerial insights about optimal scheduling of surgeries in an Outpatient Procedure Center (OPC). The scheduling rules set surgery sequences and patient ap-

pointment times to balance resource utilization and patient waiting time. The rules are tested using a discrete-event simulation based on an OPC in a large medical center. We compare the performance of easy to implement scheduling rules with more advanced optimization methods. The authors found simple rules can improve schedules used in practice. Sequencing surgeries from the longest to shortest mean duration (LPT) causes high expected overtime, and should be avoided, while sequencing from the shortest to longest mean duration (SPT) performed quite well in our experiments. The authors found that expending greater computational effort with more sophisticated optimization methods did not lead to substantial improvements. However, controlling daily procedure mix may achieve substantial improvements in performance.

Reducing Surgical Ward Congestion through Improved Surgical Scheduling and Uncapacitated Simulation
Vincent S. Chow, Martin L. Puterman, Neda Salehirad, Wenhai Huang, Derek Atkins

Frequent surgical cancellations often lead to increased stress for patients, management, and staff. A major cause of these cancellations is the unavailability of downstream beds for post-surgical recuperation. Since a portion of surgical patients are emergency patients who must be treated immediately, the only managerial lever besides adding capacity is to improve the scheduling of elective surgeries. Our paper suggests that by using a mixed-integer programming model to construct new elective surgical schedules, the frequency of surgical cancellations as a result of bed shortages can be reduced without decreasing the volume of elective surgeries. Practical guidelines derived from model solutions suggest that surgeons requiring the same downstream ward beds should be scheduled together and surgeons with long length of stay patients should be scheduled at the beginning and end of the week. In addition, we found that surgical planners were receptive to using an Excel based simulation tool to evaluate the impact of different schedules on surgical bed occupancy. This tool doubles as a valuable learning aid to understand surgical patient flow. To facilitate the implementation of these results and other advanced analytical methodologies, managers should strive to better integrate operational level data into health information systems.

Reducing Boarding in a Post-Anesthesia Care Unit
Carter Price, Bruce Golden, Michael Harrington, Ramon Konewko, Edward Wasil, William Herring

After surgery, patients recover in a post-anesthesia care unit (PACU) and are then moved to an intensive care unit (ICU). If beds are not available in the ICU, then patients stay overnight (they are boarded) in the PACU. In practice, PACU boarding is not desirable. Unplanned, overnight shifts to cover the PACU must be arranged, usually at the last minute, and this increases the stress on the hospital staff. In addition, there is evidence that PACU boarding extends a patient's length of stay and increases costs. In this paper, we develop mathematical models that are aimed at improving the flow of patients through the surgical system. These models provide hospital administrators with a flexible approach to balancing surgical admissions and discharges that can greatly reduce the number of boarders.

The Effect of Integrated Scheduling and Capacity Policies on Clinical Efficiency
Denise L. White, Craig M. Froehle, and Kenneth J. Klassen

After studying a busy outpatient clinic environment and simulating a variety of scenarios, this research yielded several useful managerial insights. First, scheduling low-variance patient visits earlier in the clinic, and high-variance visits later, slightly improves operational performance by (a) reducing patient waiting, (b) increasing physician utilization, and (c) reducing overall clinic duration. However, the effect of choosing a poor scheduling policy is not markedly worse than having no policy at all. Second, there are no significant operational advantages to adding exam rooms beyond the point where they no longer act as a bottleneck. We observed that having more exam rooms than is minimally necessary simply shifts patient waiting from the waiting room to the exam room with no net change in total patient waiting or physician utilization. Third, when keeping medical staff busy is a concern, far more attention should be paid to having the correct number of exam rooms than to patient scheduling policies, as it was far more influential on physician utilization. Finally, even though a simple patient visit process was analyzed, including ancillary services (such as radiology) did not change the recommended policies. This suggests that even complex care environments may benefit from adhering to these principles.

Reserving capacity for urgent patients in primary care
Gregory Dobson, Sameer Hasija, Edieal J. Pinker

In this paper, we investigate the trade-off between a policy of reserving certain capacity for urgent patients in primary care and a policy of no reservation also known as advanced-access. Our main result shows that in a medical practice where supply and demand are well matched, the optimal policy not only depends

Production and Operations Management 20(3), pp. ix–xii, © 2011 Production and Operations Management Society

on the relative cost of handling an urgent patient during overtime and the cost of delaying a patient, but also on the arrival dynamics of the urgent and routine patients. We find that the advanced-access policy is optimal when either the relative cost of handling an urgent patient during overtime is low compared to the cost of delaying a patient, or if a high fraction of routine patients call-in for same-day appointment, giving more urgent patients a chance to be scheduled during regular hours on a particular day. This observation supports the assertions of advocates for advanced-access who suggest that physicians who plan to move their patient scheduling policy to one similar to the advanced-access policy should encourage their patients to call for an appointment on the same day that they want to be seen by the physician.

An Analytical Framework for Designing Community-Based Care for Chronic Diseases
Beste Kucukyazici, Vedat Verter, Nancy E. Mayo

In this study, we propose a methodological framework that focuses on the community-based delivery processes for chronic diseases. While developing community-based care re-design programs for a specific chronic disease, one key question is whether the health outcomes are merely natural consequences of age, gender, co-morbidities etc. or can they be also caused by inadequacies in the care delivery processes? If the different patterns of care independently predict the health outcomes, then it would be possible to design interventions to improve the system. In answering the question above, we develop a framework to assess the care-provider visits and identify the patterns of care. These patterns of care can provide useful information to facilitate better understanding of the existing processes and to help identify possible problems with accessibility to care. Consequently, we estimate the extent to which patterns of care independently impact the health outcomes. We also estimate the potential influence of various facilitated care policies on health outcomes. A solid understanding of the potential impact of interventions would be helpful not only for the policy makers in designing community-based care delivery systems and allocating limited healthcare resources, but also for clinicians in caring for patients and researchers in designing clinical trials.