

Optimizing Intensive Care Unit Discharge Decisions with Patient Readmissions

Carri W. Chan

Division of Decision, Risk, and Operations, Columbia Business School, New York, New York 10027, cwchan@columbia.edu

Vivek F. Farias

Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, vivekf@mit.edu

Nicholas Bambos

Department of Electrical Engineering and Department of Management Science and Engineering, Stanford University, Stanford, California 94305, bambos@stanford.edu

Gabriel J. Escobar

Kaiser Permanente Division of Research, Oakland, California 94612, gabriel.escobar@kp.org

This work examines the impact of discharge decisions under uncertainty in a capacity-constrained high-risk setting: the intensive care unit (ICU). New arrivals to an ICU are typically very high-priority patients and, should the ICU be full upon their arrival, discharging a patient currently residing in the ICU may be required to accommodate a newly admitted patient. Patients so discharged risk physiologic deterioration, which might ultimately require readmission; models of these risks are currently unavailable to providers. These readmissions in turn impose an additional load on the capacity-limited ICU resources.

We study the impact of several different ICU discharge strategies on patient mortality and total readmission load. We focus on discharge rules that prioritize patients based on some measure of criticality assuming the availability of a model of readmission risk. We use empirical data from over 5,000 actual ICU patient flows to calibrate our model. The empirical study suggests that a predictive model of the readmission risks associated with discharge decisions, in tandem with simple index policies of the type proposed, can provide very meaningful throughput gains in actual ICUs while at the same time maintaining, or even improving upon, mortality rates. We explicitly provide a discharge policy that accomplishes this. In addition to our empirical work, we conduct a rigorous performance analysis for the family of discharge policies we consider. We show that our policy is optimal in certain regimes, and is otherwise guaranteed to incur readmission related costs no larger than a factor of $(\hat{\rho} + 1)$ of an optimal discharge strategy, where $\hat{\rho}$ is a certain natural measure of system utilization.

Subject classifications: dynamic programming; healthcare; approximation algorithms.

Area of review: Policy Modeling and Public Sector OR.

History: Received October 2009; revisions received May 2010, December 2010, August 2011, February 2012, April 2012; accepted July 2012.

1. Introduction

The intensive care unit (ICU) is the designated location for the care of the sickest and most unstable patients in a given hospital. These units are among the most richly staffed in the hospital: for example, in California, licensed ICUs must maintain a minimum nurse-to-patient ratio of one-to-two. Critically ill patients, who may be admitted to a hospital because of multiple illnesses, including trauma, need urgent admission to the ICU. Although it is possible to hold these patients in other areas (e.g., the emergency department) pending bed availability, this is quite undesirable, because delays in providing intensive care are associated with worse outcomes (Chalfin et al. 2007). Consequently, in such situations, clinicians may elect to discharge a patient currently in the ICU to make room for a more acute patient. For the sake of precision, we will refer to this as a demand-driven

discharge. In theory, the patient selected for such discharge would be one who was sufficiently stable to be transferred to a less richly staffed setting (such as the transitional care unit (TCU) or medical surgical floor), and, ideally, the term “stable” would be one based on ample clinical data. In practice, because predictive models of patient dynamics are not readily available, clinicians must make these transfer decisions based entirely on clinical judgment. It is natural to conjecture that demand-driven discharges might be associated with *costs*; namely, the following:

- *Patient Health Related Costs.* Patients subject to a demand-driven discharge could potentially face additional risks of physiological deterioration. Such deterioration might ultimately require readmission. Even worse, readmitted patients tend to require longer stays in the ICU and have a higher mortality rate than first-time patients (see Snow et al. 1985, Durbin and Kopel 1993).

- *System Related Costs.* Readmitted patients impose an additional load on capacity-limited ICU resources. Ultimately this hampers access to the ICU for *other* patients.

Thus motivated, the present work examines the potential benefits of a quantitative decision support system for clinicians when faced with the requirement to identify a patient for discharge in order to make room for a more acute patient. The hope is that the availability of such a system could lead to both better patient outcomes and *simultaneously* increase efficiencies in the use of scarce ICU resources. More formally, associating a demand-driven discharge with some cost that depends on the physiological characteristics of the patient discharged, our goal is to “optimally” discharge patients so as to minimize total expected costs associated with demand-driven discharges over time. One example of such a cost may be the increase in mortality risk because of a demand-driven discharge. As a second example, one might consider the increase in expected readmission load associated with the increased likelihood of readmission because of a demand-driven discharge. We will eventually estimate and test several such cost metrics.

Our analysis will consider a stylized model of an actual ICU where the number of ICU beds is fixed.¹ Patients arrive to the ICU at random times; patients are categorized into a finite number of classes based on their physiological characteristics upon admission. There exist a number of proprietary classification systems based on a patient’s physiological characteristics. All new arrivals must be given an ICU bed immediately; they cannot queue up and wait for a bed to become available. This models the aforementioned fact that new ICU patients are typically extremely high priority. If no beds are vacant upon the arrival of a new patient, a current patient will have to be discharged in order to accommodate the newly arriving patient.² The demand-driven discharge of a patient will incur a cost that depends on that patient’s class; this cost is modeled to reflect the impact of the demand-driven discharge on the patient as well as the system as described above. Our goal will be to minimize the expected costs incurred because of demand-driven discharges over some finite horizon. This is a difficult problem, and our analysis of this stylized model will suggest simple policies for which we will develop performance guarantees. More interestingly, we will conduct a detailed simulation study based on real data to examine our recommendations.

1.1. Our Contributions

We make the following key contributions:

- *Interpretability.* We show that a *myopic* policy is a potentially good approximation to an optimal policy. This corresponds to an index policy wherein every patient class is associated with a class specific index. The index for a given class can be computed from historical patient flow data in a robust fashion. Depending on the cost metric under consideration, we will demonstrate that these indices

can serve as natural measures for patient criticality that have both clinical as well as operational merit. The index policy then has an appealing clinical interpretation: when a patient must be discharged in order to accommodate new patients, one simply discharges an existing patient of the lowest possible criticality index.

- *Robustness.* Our index policy is “robust”. In particular the indices we compute are oblivious to patient traffic intensities that are highly variable and difficult to estimate. Rather, they rely on quantities relevant to specific classes of patients that are typically far simpler to estimate from data. For the data set under consideration, relative changes of estimated parameters greater than 50% were typically required to induce a change in the associated indices.

- *Performance Guarantees and Operational Relevance.* We demonstrate via a theoretical analysis that our index policy is, for a certain class of problems, optimal and in general incurs total expected cost that is no more than $1 + \hat{\rho}$ times that incurred under an optimal discharge rule, where $\hat{\rho}$ is a certain natural measure of ICU utilization. We identify a cost metric—the increase in expected readmission load because of a demand-driven discharge—that in addition to enjoying a clinical interpretation as a measure of criticality, can be shown to capture a notion of throughput optimality.

- *Empirical Validation.* Most importantly, we calibrate our model to empirical data from over 5,000 patient flows at a large privately owned partnership of hospitals and identify parameters for patient dynamics. We consider a variety of cost metrics, including several natural metrics motivated by existing clinical literature and modifications of these cost metrics such as the operationally relevant metric alluded to above. We measure the impact of these discharge policies along two dimensions. First, to understand impact at the individual patient level, we measure mortality rates under the various policies. Second, to understand system level impact we measure the readmission load incurred under the various policies. In doing so, we identify a policy that, in addition to fitting within the ethos of ordering patients by a measure of criticality, has substantive benefits over other, perhaps more “obvious” policies: Under modest assumptions on patient traffic, it incurs a 30% reduction in readmission load at no cost to mortality rate.

As such, this study provides a framework for the design of demand-driven discharge policies and in doing so identifies a policy that allows us to utilize available ICU resources as effectively as possible while not sacrificing the quality of patient outcomes. At a high level, our analysis suggests that investments in providing clinicians with more decision support (e.g., severity of illness scores and the associated risks of physiological deterioration) could translate into tangible benefits both in terms of improved patient outcomes, increased efficiency, and decreased costs.

1.2. Related Literature

The use of critical care is increasing, which is making already limited resources even more scarce (Halpern and

Pastores 2010). In fact, it was shown that 90% of ICUs will not have the capacity to provide beds when needed (Green 2003). As such, it is the case that some patients may require premature discharges in order to accommodate new, more critical patients. In a recent econometric study (Kc and Terwiesch 2012), these types of patient discharges were shown to be a legitimate cause of patient readmissions thereby effectively reducing peak ICU capacity because of the additional load the readmitted patients bring. The empirical data we have analyzed in calibrating our ICU model corroborates this fact.

There has been a significant body of research in the medical literature that has looked at the effects of patient readmissions. In Chrusch et al. (2009), high occupancy levels were shown to increase the rate of readmission and the risk of death. Unfortunately, readmitted patients typically have higher mortality rates and longer hospital lengths of stay (see Franklin and Jackson 1983, Chen et al. 1998, Chalfin 2005, Durbin and Kopel 1993 and related works).

When a new patient arrives to the ICU, either after experiencing some trauma or completing surgery, he must be admitted. If there are not enough beds available, space must be allocated by transferring current patients to units with lower levels of staffing and care. In Swenson (1992) and related works, the authors examine how to allocate ICU beds from a qualitative perspective that is not based on analysis of patient data but rather on philosophical notions of “fairness.” The authors propose a five-class ranking system for patients based on the amount of care required by the patient as well as his risk of complications. Our approach may be seen as a quantitative perspective on the same problem wherein decisions are motivated by the analysis of relevant quantitative patient data. To date, the work (particularly in the medical community) on how to determine discharge decisions has been rather subjective because of the lack of information-rich models that attempt to capture patient dynamics. Thus, these works (see for instance Bone et al. 1993 and a study by the American Thoracic Society 1997) have not considered that discharging a patient from the ICU in order to accommodate new patients may result in readmission, further increasing demand for the limited number of beds and ultimately compromising the quality of care for all patients involved. We not only propose such a model, but also show the efficacy of discharge policies that utilize this previously unavailable information.

Dobson et al. (2010) consider a setup quite similar to ours but ignore the readmission phenomenon; rather they simply seek to quantify the total expected number of patients discharged in order to accommodate new, more critical patients. To this end, they analyze a policy that chooses to discharge patients with the shortest remaining service time (which are modeled as deterministic quantities). As will be seen in §5, which presents an empirical performance evaluation using a real patient flow data set,

a distinct heuristic is desirable when one does account for patient readmission.

A number of modeling approaches have been used to make capacity, staffing, and other tactical decisions in the healthcare arena (see for instance Huang 1995, Kwak and Lee 1997, Green et al. 2006a). Queueing theory has been particularly useful to study the question of necessary staffing levels in hospitals. As examples of this work, Green et al. (2006b) and Yankovic and Green (2011) consider a number of staffing decisions from a queueing perspective. The goal is to provide patients with a particular service level (in terms of timeliness, and also nurse-to-patient ratio) while at the same time addressing issues such as temporal variations in arrival rates of patients of different types. See also Green (2006) for an overview of the use of operations research models for capacity planning in hospitals. Murray et al. (2007) consider different factor such as age, gender, physician availability, and number of visits per patient per year to determine the largest patient panel size that may be supported by available resources. In Green and Savin (2008), the authors consider how to reduce delay in primary care settings by varying the number of patients served by the particular primary care office. When a patient wishes to make an appointment, he may be delayed before the physician is able to see him. Two significant differences separate the problem we consider from those considered in the above streams of work: arriving patients to an ICU must receive service immediately (which thus necessitates discharging current patients). This in turn requires that we consider individual patient dynamics and in particular model the impact of discharging a patient to accommodate new ones on the discharged patient’s likelihood of revisiting the ICU. We can then make staffing decisions in much the same way as the aforementioned work.

In a related paper on ICU patient flow (Shmueli et al. 2003), the authors examine the affect of ICU admission strategies on the distribution of ICU bed occupancy. The authors assume it is possible for patients to wait for an ICU bed, regardless of their criticality. For the specific ICUs we consider, waiting is highly undesirable (thereby necessitating our modeling decisions that arriving patients be given a bed immediately). An interesting direction for future work would be to consider an intermediate scenario, where some patients may be delayed, whereas others must be given a bed immediately.

Finally, relative to recent work by Chan and Farias (2009), we note that the present paper considers a class of models entirely distinct from the “depletion problems” studied there and succeeds in establishing relative approximation guarantees for a class of models left unaddressed by that past work. The properties we exploit in our analysis are new and it would be interesting to understand whether the techniques introduced here have application to the more natural cost-minimization variants of the queueing problems introduced in Chan and Farias (2009).

The rest of the paper proceeds as follows. Section 2 formally introduces the queueing model and patient dynamics we study. In §3, we analyze the performance of an index policy that selects patients to discharge in a greedy manner based on their expected costs incurred because of demand-driven discharges. We explore a scenario where the proposed greedy policy (based on an information-rich model) is, in fact, optimal. Furthermore, in a more general setting, we show that the greedy policy is guaranteed to be within a factor of $(\hat{\rho} + 1)$ of optimal, where $\hat{\rho}$ is a measure of system utilization. In §4, we discuss various measures of criticality that constitute clinically relevant cost metrics. These measures include an important refinement to a criticality measure that has received some attention in the critical care literature. In §5, we discuss the calibration of our model using a proprietary ICU patient flow data set from a group of private hospitals. Having calibrated our model, we show in §6 that our primary proposal outperforms a number of benchmarks of interest. We conclude in §7.

2. Model

We begin by proposing a stylized model of the patient flow dynamics in a hospital ICU and account for the fact that discharging a current ICU patient in order to accommodate a new one is undesirable for the discharged patient and comes at a “cost.” At a high level, our model captures the fact that a newly admitted patient must receive ICU resources and that this requirement in turn could necessitate the discharge of an existing ICU patient. Such a discharged patient may suffer physiologic deterioration because of the demand-driven discharge. Because arriving patients cannot be queued or blocked, the model we consider is distinct from a typical queueing model. Presuming a measure of cost associated with a demand-driven discharged patient, a natural goal is to find a patient discharge policy that minimizes this cost.

Preliminaries. We consider time to be discrete and indexed by $t \in [0, T]$. In each time slot, we must determine if a patient must be discharged and, if so, which one. If there are enough available beds to accommodate all current and arriving patients, discharge of current patients is not required.

We assume that patients may be classified into one of M classes, each potentially corresponding to the particular ailment/health condition of the ICU patient. Let $m \in \mathcal{M} = \{1, 2, \dots, M\}$ denote the type of a particular patient. Patients from a given class are assumed to have identical statistics for their initial lengths of stay and identical costs associated with a demand-driven discharge. Specifically, we assume that the initial length of stay for a patient of class m is a geometric random variable with mean $1/\mu_m^0$. If such a patient is discharged prior to completing treatment because of the arrival of a more acute patient, a cost, $\phi_m \geq 0$, is incurred. Although the patient length of stay distribution is assumed to be memoryless for the purposes of analysis,

our empirical study assumes log-normal distributions for length of stay that are fit to the empirical data (see §5). Finally, in §3.3, we discuss an extension to our model that is able to capture a patient’s evolution and changing condition during his ICU stay by using a phase-type length-of-stay distribution.

At most one new patient can arrive in each time slot and an arrival occurs with probability λ . We define $\hat{\rho} = \lambda / \min_m \mu_m^0$ as a measure of the utilization of the ICU: a higher $\hat{\rho}$ implies a more stressed ICU and a lower value implies more able bed resources. Notice that this measure does not rely on the relative arrival intensities of various patient types. We let $a_{t,m}$ denote the probability that a newly arriving patient at time t is of type m . These probabilities are deterministic and known a priori to the optimal discharge policy; the policy we study will require neither knowledge of λ nor the probabilities $a_{t,m}$.

We assume that the ICU has B beds. If all B beds are full and a new patient arrives, then a patient must be discharged prior to completing service in order to accommodate the newly arrived patient. We let $x_{t,m} \in \{0, 1, \dots, B\}$ denote the number of class m patients currently in the ICU at the beginning of time slot t and let $y_{t,m} \in \{0, 1\}$ be an indicator for the arrival of a type m patient at the start of the t th epoch. Note that because at most one new patient can arrive in each time slot, $\sum_{m=1}^M y_{t,m} \leq 1$ for all t . A current patient must be discharged if $\sum_{m=1}^M x_{t,m} + \sum_{m=1}^M y_{t,m} = B + 1$; we refer to this type of discharge as a demand-driven discharge. The natural departure (or service completion) of patient type m occurs at the end of the t th time slot with probability μ_m^0 after any demand-driven discharge and/or admission occurs.

State and action space. The dynamic optimization problem we will propose is conveniently studied in a “state space” model. We define our state-space as the set

$$\mathcal{S} = \left\{ (x, y, t) : x \in \{0, 1, \dots, B\}^M, \sum_{m=1}^M x_m \leq B, y \in \{0, 1\}^M, \sum_{m=1}^M y_m \leq 1, 0 \leq t \leq T \right\}.$$

In particular, the state of the system is completely described by the number of patients of each type currently in the ICU, the type of the arriving patient at that state if any, and the epoch in question. We denote by $x(s)$ the projection of s onto its first coordinate and similarly employ the notation $y(s)$ and $t(s)$. We let the random variable $s_t \in \mathcal{S}$ denote the state in the t th epoch. Note that because the $\{a_{t,m}\}$ process is assumed to be deterministic and given a priori, the current time slot t completely specifies the arrival probabilities for each patient class.

For each state s , let $\mathcal{A}(s) \subset \mathcal{M}$ denote the set of feasible actions that can be taken in time slot $t(s)$. For states wherein a demand-driven discharge is required, i.e., states s for which $\sum_m x(s)_m + y(s)_m > B$, we have

$\mathcal{A}(s) = \{m: x(s)_m > 0\}$. At all other states s , $\mathcal{A}(s) = \{m: x(s)_m > 0\} \cup \{0\}$. Thus, an action $A \in \mathcal{A}(s)$ specifies the class of the patient, if any, to be discharged in time slot $t(s)$; because only one patient can arrive in each time slot, at most one demand-driven patient discharge is required to accommodate a new patient. We will henceforth suppress the dependency of the set of feasible actions, $\mathcal{A}(s)$, on s .

Dynamics. Let $s' = S(s, A)$ denote the random next state encountered upon employing action A (demand-driven discharge of patient type A) in state s . A random number, $X_{t(s),m}$, of class m patients will complete treatment and depart naturally, where $X_{t(s),m}$ is a binomial- $(x(s)_m + y(s)_m - \mathbf{1}_{\{A=m\}}, \mu_m^0)$ random variable. Hence, R_t be independent random variables, defined for each t , indicating the type of an arriving patient at the start of the t th epoch. Hence, R_t takes values in $\{1, 2, \dots, M\} \cup \{0\}$; $R_t = m$ with probability $\lambda a_{t,m}$ for $m \in \{1, 2, \dots, M\}$ and $R_t = 0$ with the remaining probability. The vector denoting arrivals at the next state, $Y_{t(s)+1}$ is then given by $Y_{t(s)+1, m} = \mathbf{1}_{\{R_{t(s)+1}=m\}}$. Thus, $s' = S(s, A)$ is defined as

$$x(s')_m = x(s)_m + y(s)_m - \mathbf{1}_{\{A=m\}} - X_{t(s),m},$$

$$y(s')_m = Y_{t(s)+1, m},$$

$$t(s') = t(s) + 1.$$

Cost function. The cost incurred for taking action A is defined by a cost function $C: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}_+$. Such a cost function might capture a number of quality metrics. For instance, the cost function might reflect the net decrease in quality-adjusted life years (QALYs) as a result of a demand-driven discharge. Our discussion is able to capture any such cost function. We take $C(s, A) = \phi_A$ for $A \in \{1, 2, \dots, M\}$, and $C(s, 0) = 0$. In §4, we discuss clinically relevant cost metrics.

Objective. Let Π denote the set of feasible discharge policies, π which map the state space \mathcal{S} to the set of feasible actions \mathcal{A} . Define the expected total cost-to-go under policy π as

$$J^\pi(s) = E \left[\sum_{t'=t(s)}^{T-1} C(s_{t'}, \pi(s_{t'})) \mid s_{t(s)} = s \right].$$

We let $J^*(s) = \min_{\pi \in \Pi} J^\pi(s)$ denote the minimum expected total cost-to-go under any policy. We denote by π^* a corresponding optimal policy, i.e., $\pi^*(s) \in \arg \min_{\pi \in \Pi} J^\pi(s)$.

The optimal cost-to-go function (or *value function*) J^* and the optimal discharge policy π^* can in principle be computed numerically via dynamic programming; in particular, define the dynamic programming operator \mathcal{H} according to

$$(\mathcal{H}J)(s) = \min_{A \in \mathcal{A}} E[C(s, A) + J(S(s, A))] \quad (1)$$

for all $s \in \mathcal{S}$ with $t(s) \leq T - 1$. The optimal value function J^* may then be found as the solution to the Bellman

equation $\mathcal{H}J = J$, with the boundary condition $J(s') = 0$ for all s' with $t(s') = T$. The optimal policy π^* may be found as the greedy minimizer with respect to J^* in (1). The minimization takes into consideration the current state s , the distribution of future patient arrivals, as well as the impact of the current decision on future states. References to an optimal policy in subsequent sections will refer to precisely this policy. The size of \mathcal{S} precludes this straight-forward dynamic programming approach. Moreover, even if an optimal solution were possible, the robustness of such an approach and its implementability remain in question because it relies on detailed patient arrival statistics that are typically not stationary and difficult to estimate. As such, our goal will be to design simple, robust heuristics for the load minimization problem at hand.

In addition to the above objective, one may also consider the task of finding an *average-cost* optimal policy; i.e., the task of finding a *stationary* policy π (a policy that satisfies $\pi(s) = \pi(s')$ for all s, s' with $x(s) = x(s')$, and $y(s) = y(s')$), that solves

$$\kappa^*(s) = \min_{\pi} \kappa^\pi(s),$$

where $\kappa^\pi(s) = \limsup_{T \rightarrow \infty} (1/T) E[\sum_{t'=t(s)}^{T-1} C(s_{t'}, \pi(s_{t'})) \mid s_{t(s)} = s]$ is the average-cost to go (i.e., the long-run costs incurred because of demand-driven discharges) under policy π .

It is not difficult to see that the Markov chain on $\hat{\mathcal{S}}$ (the projection of \mathcal{S} on its x and y coordinates) induced under any stationary policy π is irreducible, so that in fact, the above problem is solved simultaneously for all s by a common stationary policy π^* , and $\kappa^\pi(s) = \kappa^\pi$ for all $s \in \mathcal{S}$ and a stationary policy π . Finally, the ergodic theorem for Markov chains implies (with some abuse of notation), that

$$\kappa^\pi = \sum_{s \in \hat{\mathcal{S}}} \nu_\pi(s) C(s, \pi(s)),$$

where ν_π is the stationary distribution induced by π on $\hat{\mathcal{S}}$.

3. A Priority-Based Policy

This section introduces an index policy for the dynamic optimization problem proposed. Under such a policy, the patient selected for a demand-driven discharge is simply chosen from a patient class that would incur the minimal cost. In particular, such a policy states that the patient (class) $\pi^g(s)$ chosen for discharge satisfies

$$\pi^g(s) \in \arg \min_{A \in \mathcal{A}(s)} C(s, A) = \arg \min_{m \in \mathcal{A}(s)} \phi_m. \quad (2)$$

It is easy to see that the policy specified by (2) has a natural implementation as an “index” policy. It is interesting to note that implementing such a policy requires data about particular patient classes but does not require the estimation of arrival rates of the various classes. This latter information is highly dynamic and difficult to estimate.

Because the policy we have proposed ignores the effect of future arrivals and the expected length of stay of the current occupants, it is natural to expect such a policy to be suboptimal. In the appendix, example A shows what can go wrong.

In light of the suboptimality of our proposed priority based policy, the remainder of this section is devoted to establishing performance guarantees for this policy. In particular, we identify a setting where the greedy policy is, in fact, optimal. More generally we establish that the greedy policy incurs expected costs that are at most a factor of $(\hat{\rho} + 1)$ times the expected costs incurred by an optimal policy (i.e., the greedy policy is a “ $(\hat{\rho} + 1)$ -approximation”), where $\hat{\rho} = \lambda / \mu_{\min}^0$ (here $\mu_{\min}^0 \triangleq \min_m \mu_m^0$) is a measure of the utilization of the ICU defined in §2: a higher $\hat{\rho}$ implies a more stressed ICU and a lower value implies more able bed resources. This latter bound is independent of all other system parameters.

3.1. Greedy Optimality

In this section, we consider a special case of the general model presented in §2 for which a greedy discharge rule is optimal. The proof of this result can be found in the appendix. In particular we have the following theorem:

THEOREM 1 (GREEDY OPTIMALITY). *Assume that for any two patient classes i, j with $\phi_i \leq \phi_j$ we also have $1/\mu_i^0 \geq 1/\mu_j^0$. Then, we have that the greedy policy is optimal, i.e.,*

$$J^g(s) = J^*(s), \quad \forall s \in \mathcal{S}.$$

Theorem 1 considers problems for which patients with lower cost also have higher nominal lengths of stay. In this case, because eliminating a low-cost patient also frees up capacity that would have otherwise been occupied for a relatively longer time, it is intuitive to expect the greedy policy to be optimal. However, the assumptions of the theorem are likely to be restrictive in practice. In the next section, we consider the performance of the greedy policy without any assumptions on problem primitives.

3.2. A General Performance Guarantee

Our objective in this section is to demonstrate that the greedy heuristic incurs expected costs that are within $\hat{\rho} + 1$ times that incurred by an optimal policy as discussed in §2. In particular, we will show that for any state $s \in \mathcal{S}$, $J^g(s) \leq (\hat{\rho} + 1)J^*(s)$, where $\hat{\rho} = \lambda / \mu_{\min}^0$ is a utilization ratio defined in §2.

To show the desired bound, we begin with a few preliminary results for the optimal value function J^* . The proofs of these results can be found in the appendix. The first result is a natural monotonicity result that says that having an ICU with higher occupancy levels is less desirable than having lower occupancy levels. In particular,

LEMMA 1 (VALUE FUNCTION MONOTONICITY). *For all states $s, s' \in \mathcal{S}$ satisfying $x(s) \geq x(s')$, $y(s) = y(s')$, $t(s) = t(s')$, we have*

$$J^*(s) \geq J^*(s').$$

In words, the Lemma 1 states that all else being equal, it is advantageous to start at a state with a fewer number of patients occupying the ICU. Now suppose in state s we chose to take the greedy action as opposed to the optimal action (assuming of course that the two are distinct). It must be that the former leads to a higher cost state than does the optimal action. The following result places a bound on this cost increase. In particular, we have

LEMMA 2 (ONE STEP SUBOPTIMALITY). *For any state $s \in \mathcal{S}$ and $\alpha = \hat{\rho} / (\hat{\rho} + 1)$,*

$$E[J^*(S(s, \pi^g(s)))] \leq \alpha C(s, \pi^*(s)) + E[J^*(S(s, \pi^*(s)))].$$

In words, Lemma 2 tells us that if we were to deviate from the optimal policy for a *single* epoch (say, in state s), the impact on long-term costs is bounded by the quantity $\alpha C(s, \pi^*(s))$. We now use this bound on the cost of a single period deviation in an inductive proof to establish performance loss incurred in using the greedy policy; we show that the greedy heuristic is guaranteed to be within a factor of $\hat{\rho} + 1$ of optimal, where $\hat{\rho} = \lambda / \mu_{\min}^0$ is the utilization ratio of the ICU defined in §2.

THEOREM 2. *For all $s \in \mathcal{S}$, $J^g(s) \leq (\hat{\rho} + 1)J^*(s)$.*

PROOF. The proof proceeds by induction on the number of time steps that remain in the horizon, $T - t(s)$. The claim is trivially true if $t(s) = T - 1$ because both the myopic and optimal policies coincide in this case. Consider a state s with $t(s) < T - 1$ and assume the claim true for all states s' with $t(s') > t(s)$.

Now if $\pi^*(s) = \pi^g(s)$ then the next states encountered in both systems are identically distributed so that the induction hypothesis immediately yields the result for state s . Consider the case where $\pi^*(s) \neq \pi^g(s)$. Defining $\alpha = \hat{\rho} / (\hat{\rho} + 1)$, we have

$$\begin{aligned} J^*(s) &= C(s, \pi^*(s)) + E[J^*(S(s, \pi^*(s)))] \\ &\geq (1 - \alpha)C(s, \pi^*(s)) + E[J^*(S(s, \pi^g(s)))] \\ &\geq (1 - \alpha)C(s, \pi^g(s)) + E[J^*(S(s, \pi^g(s)))] \\ &\geq (1 - \alpha)C(s, \pi^g(s)) + E[(1 - \alpha)J^g(S(s, \pi^g(s)))] \\ &= (1 - \alpha)J^g(s) \\ &= \frac{1}{\hat{\rho} + 1}J^g(s). \end{aligned} \tag{3}$$

The first equality comes from the definition of the optimal policy. The first inequality comes from Lemma 2. The second inequality comes from the definition of the greedy policy that minimizes single period costs. The third inequality

comes from the induction hypothesis. The second equality comes from the definition of the greedy value function. This concludes the proof.

Our guarantee on performance loss suggests that in regimes where ICU utilization is low, the greedy policy is guaranteed to be close to optimal. At some level, this is an intuitive result—low levels of utilization should imply infrequent demand-driven discharges as there are likely to be available beds when new patients arrive; Theorem 2 makes this intuition precise by demonstrating a bound on how performance loss scales with utilization levels. Our guarantees are worst case; later in this section we will consider a generative family of problems for which the performance loss is a lot smaller than predicted, even at high utilization levels. Moreover, we will demonstrate via an empirical study using patient flow data, that the greedy policy is superior to a number of benchmarks that resemble current practice. Before we continue, we briefly discuss extensions to the model presented in §2 and how the presented results can be applied.

3.3. Patient Evolution During ICU Stay

Thus far, we have assumed the distribution for the length of stay of each patient is memoryless. Because the health of a patient will vary over the course of his stay, one may wish to employ a length-of-stay distribution that does not have a constant hazard rate. We now consider how to incorporate this more realistic scenario.

For each patient class m , consider a random progression of the state of their health condition. Let $h^m \in \{h_0^m, h_1^m, \dots, h_{n_m}^m\}$ denote the set of health condition states patient class m can achieve. Whenever a new patient of type m arrives, it begins with a health state of h_0^m . Assuming that a patient is in health state h_n^m in some epoch, the patient departs with probability $\mu_m^0(h_n^m)$. If he does not depart, he evolves to health state h_{n+1}^m with probability γ_n^m and remains in state h_n^m with probability $1 - \gamma_n^m$. Should a patient in health state h_n^m be demand-driven discharged, the cost he introduces is $\phi_m(h_n^m)$. The different health condition states and corresponding departure probabilities enable us to capture the changes (improvement or deterioration) in patient health as a patient spends time in the ICU. Note that there are no constraints on the relationship between the $\mu_m^0(h_n^m)$ so that the patient does not necessarily improve with time. Indeed, there have been studies that show that patients likelihood of departure *decreases* the longer they have spent in the hospital (Chalfin 2005).

The state space now needs to be expanded to incorporate the different health states each patient class can achieve. To do this, we can redefine $x(s)$ to be a two-dimensional array, where $x_{m,n}(s)$ denotes the number of class m patients in health condition state h_n^m . We consider using the natural analogue to the greedy policy discussed thus far:

$$\tilde{\pi}^s(s) \in \arg \min_{(m,n): x_{m,n}(s) > 0} \phi_m(h_n^m).$$

Now, Lemma 1 can be established exactly as before for this new system, with the understanding that we will say $x(s) \geq x(s')$ iff $x_{m,n}(s) \geq x_{m,n}(s')$ for all m, n . Further, the analysis used in the proof of Lemma 2 also applies identically as in the case of that result to show that for $\alpha = \tilde{\rho}/(\tilde{\rho} + 1)$,

$$E[J^*(S(s, \tilde{\pi}^s(s)))] \leq \alpha C(s, \pi^*(s)) + E[J^*(S(s, \pi^*(s)))],$$

where we now define

$$\tilde{\rho} = \frac{\lambda}{\min_{m,n} \mu_m^0(h_n^m)}.$$

With these results, the proof of Theorem 2 applies verbatim to yield the following:

THEOREM 3. For all $s \in \mathcal{S}$, $J^{\tilde{\pi}^s}(s) \leq (\tilde{\rho} + 1)J^*(s)$.

3.4. Patient Diversions

Throughout our discussion we have assumed that all new patients *must* be given a bed immediately. In some cases, high occupancy levels in an ICU can lead to congestion in other areas of the hospitals, such as the Emergency Department (ED), because patients cannot be transferred across hospitals units. In Allon et al. (2012) and McConnell et al. (2005), it is shown that when ICU occupancy levels are high, ambulance diversions increase. Because of the inability to move patients from the ED to ICU, patients are blocked from the ED and ambulances must be diverted to other hospitals. In de Bruin et al. (2007), the authors examine the case of bed allocation given a maximum allowable number of patient diversions in the case of cardiac intensive care units. The authors identify scenarios where achieving the target number of patient diversions is possible, but they do not consider how to make admission and discharge decisions. Ambulance diversion comes at a cost—for both the hospital and patient. The hospital loses the revenue generated for treatment (McConnell et al. 2006, Melnick et al. 2004, Merrill and Elixhauser 2005), and delays because of transportation time may result in worse outcomes for the diverted patient (Schull et al. 2004). On the other hand, diversions can sometimes alleviate overcrowding (Scheulen et al. 2001).

Typically, diverted ambulance patients are not the ones who require ICU care (Scheulen et al. 2001). However, within a hospital it may still be possible to block new ICU patients admissions, either by diverting them to another unit (i.e., a TCU or general floor) within the same hospital or transferring them to an ICU in a different hospital (because of the integrated nature of the hospital system we study, such intrahospital transfers do occur). Blocking new patients may reduce the number of demand-driven discharges. Note that these new patients are often being transferred from a different hospital unit (emergency department, operation room, general ward, etc.) rather than being brought in by ambulances, which is the case of the

extensive body of literature on ambulance diversions. Given the ability to divert patients, we consider how to incorporate patient diversions into our model and decision analysis. We extend our model to allow new ICU patients to be diverted to another hospital ICU or unit of lesser care. Hence, when an ICU is full the hospital administrator must decide whether to block the new patient or to make a demand-driven discharge of a current patient in order to admit the new patient.

To formalize the above decision making, we consider the following extension of our model: in a given state s , we permit an additional action corresponding to diversion, which we denote by D ; we let $C(s, D)$ denote the cost associated with a diversion in state s ; as per our discussion above, this cost must capture the increased risks to the patient being diverted in state s (i.e., the arriving patient in that state) as also potential revenue losses to the hospital. We then consider employing the following policy: for states $s \notin \hat{\mathcal{S}}_{\text{full}}$, i.e., states where the ICU has available capacity, no action is necessary. Otherwise, we follow the following diversion/discharge policy:

$$\hat{\pi}(s) = \begin{cases} \pi^g(s), & \text{if } C(s, D) \geq C(s, \pi^g(s)); \\ D, & \text{otherwise.} \end{cases}$$

Now, Lemma 1 can be established exactly as before for this new system, and the analysis used in the proof of Lemma 2 also applies identically as in the case of that result to show that for $\alpha = \hat{\rho}/(\hat{\rho} + 1)$,

$$E[J^*(S(s, \hat{\pi}(s)))] \leq \alpha C(s, \pi^*(s)) + E[J^*(S(s, \pi^*(s)))].$$

Given these properties, the proof of Theorem 2 applies verbatim to yield


$$\text{THEOREM 4. For all } s \in \mathcal{S}, J^{\hat{\pi}}(s) \leq (\hat{\rho} + 1)J^*(s).$$

3.5. Comparison to Optimal

This section is devoted to examining the performance loss of the greedy policy via numerical studies. We compare the greedy and optimal policies for a set of smaller problems for which the optimal policy is actually computable. In the following section, we examine larger problem instances calibrated to empirical data and compare the performance of the greedy policy to a number of benchmark policies.

In §3.2, we have shown that the greedy performance is an $(\hat{\rho} + 1)$ -approximation algorithm to optimal. To enable computation of the optimal policy, we consider a small scenario with $B = 10$ beds, $M = 2$ patient types, and a time horizon of 240 time slots (assuming admission and discharge decisions are made every six minutes, or 10 times an hour, this corresponds to a time horizon of 24 hours). For each data point, we fix the probability of arrival of each patient type. We consider 100 different realizations for the nominal length of stay and cost of demand-driven discharge of each patient type, which we vary uniformly at random

Figure 1. Performance of greedy policy compared to optimal for varying arrival rates.



with mean 25 hours and 2.5 units of cost, respectively. For each fixed set of parameters— $a_{i,t}$, μ_i^0 , and ϕ_i —we calculate the optimal policy using dynamic programming. We compare the average performance of this optimal policy to the performance of the greedy policy over 100 sample paths.

Figure 1 shows the ratio of the greedy performance to the optimal performance ($J^g(s)/J^*(s)$) for a range of different arrival rates. As from §2, the probability of a patient arrival is given by λ and the probability an arrival is of patient type 1 is given by a_1 . Values above one show the loss in performance because of using the greedy policy. We can see that the greedy policy performs within 3% of optimal, which is substantially superior to what the bound in §3.2 suggests. In fact, for reasonable arrival rates ($\lambda < 0.05$ means one patient arrives every two hours) the performance loss of the greedy policy is less than 1% of optimal. These differences are so small they can essentially be ignored because of possible numerical errors. The greedy policy does not require arrival rate information and is much simpler to compute than optimal. These simulation results suggest that using the greedy policy results in little performance loss while significantly reducing the computational complexity. In fact, whereas the complexity of the greedy policy grows linearly in the time horizon, T , and logarithmically in the number of patient types ($\log M$), the complexity of the optimal policy grows exponentially in a number of problem parameters despite only resulting in slightly higher performance. The simplicity and good performance of the greedy policy, which simply prioritizes different patient types, makes it desirable for real-world implementation.

4. Clinical Relevance

Our exposition thus far has treated the problem of prioritizing patients for demand-driven discharges as a purely *operational* problem. In a nutshell, we have shown that if one desires to minimize some long-run cost metric impacted by demand-driven discharge decisions, then a priority rule

that is “greedy” with respect to the cost metric serves as a reasonable and operationally viable approximation to an optimal policy.

This section considers *clinical* issues relevant to the problem at hand. In particular, the clinical viability of a discharge policy is of paramount importance. In particular, what remains to be specified are clinically relevant cost metrics and priority rules that capture factors physicians would like to account for in making discharge decisions. Certainly, the general consensus of the medical community is that patients should be discharged in order of least critical first (see, for instance, Swenson 1992). However, what determines criticality is left wide open to interpretation and is highly dependent on the experience and training of an individual physician. In fact, disagreements on which patient should be discharged arise frequently and in an effort to building a process around this critical decision, many hospitals are adopting an intensivist-managed system that makes triage decisions for all patients in the ICU (Franklin et al. 1990, Task Force of the American College of Critical Care Medicine 1999). Although such a process will remain necessarily subjective, there is a strong desire that the process be informed by quantitatively designed best-practice recommendations. In this spirit, we consider several policies that fall within the ethos of a priority rule based on measures of patient criticality that have been broached in the extant medical literature.

Mortality risk. A natural measure of patient “criticality” is mortality risk. In fact, the commonly used APACHE (acute physiology and chronic health evaluation) and SAPS (simplified acute physiology score) severity scores are based on mortality predictions for ICU patients (Zimmerman et al. 2006, Moreno et al. 2005). Although it is obvious that patients with high mortality risk are “critical” and should not be demand-driven discharged, intensivists are likely to find this measure of criticality too crude to be of value in practical scenarios. To be more precise, one typically needs to be able to distinguish among patients all with relatively low mortality risk but variedly long and complex recoveries. In addition, a metric based solely on mortality risk will fail to capture a system-wide view of the ICU and in particular, the impact a discharge decision for a given patient might have on the ability to provide timely and quality care for other patients. Specifically, such a metric fails to account for the impact a discharge decision has on ICU congestion—congestion in the ICU can result in postponing surgeries, delaying admissions, and/or rerouting patients to other units—all of which are associated with worse outcomes (Metcalf et al. 1997, Mitchell et al. 1995, Smith et al. 1995, Chalfin et al. 2007, Renaud et al. 2009, Rincon et al. 2010). As such, it is ethically important to consider factors related to congestion in making such decisions.

Readmission risk. A potential refinement on using simply mortality risk as a measure of patient criticality is accounting for readmission risk. In fact, measures related to readmission risk have been gaining attention and credibility in

the medical community motivated primarily by two factors: medical outcomes and payment structures. In terms of medical outcomes, readmitted patients have been shown to be worse off, with higher mortality and longer length of stay (Chen et al. 1998, Durbin and Kopel 1993, Rosenberg and Watts 2000). Recognizing the clinical risks associated with readmissions, many hospitals are adopting discharge strategies that account for patient readmissions (Franklin and Jackson 1983, Yoon et al. 2004). In terms of monetary incentives, readmissions can also increase costs by over 25% (Naylor et al. 2004). Acknowledging the detrimental impact of readmissions on patient outcomes and the extraordinarily high costs associated with the care of readmitted patients, the Patient Protection and Affordable Care Act (2010) requires Medicare to begin reducing readmissions in 2013. Although physiology-based probabilistic models for assisting ICU physicians in making discharge decisions are not widely available, there has been recent interest in developing risk scores to assess readmission risks, similar to what the APACHE and SAPS scores do for mortality (Gajic et al. 2008). In this spirit, one may consider several concrete metrics.

A crude metric. As a concrete measure of readmission risk, one might consider the *likelihood of readmission*. One expects that such a measure will be fairly correlated with a measure of mortality risk. At the same time, such a measure will move toward addressing some of the pitfalls of using mortality risk alone. That said, such a measure remains somewhat coarse in two regards: First, it fails to account for the actual impact of the demand-driven discharge decision itself on readmission risk; because readmissions might arise because of a multitude of other factors, this is crucial. Second, it fails to account for the diversity in complications that might occur upon a readmission.

A refinement (our proposed policy). We consider a mild refinement to the above measure of readmission risk: we consider the increase in readmission load attributable to a demand-driven discharge. Roughly speaking, we can think of this refinement as accounting not only for readmissions, but in addition, the typical length of stay upon such a readmission. More precisely, let p_m^N and $1/\mu_m^{R,N}$ be the probability of readmission and expected readmission LOS of patient class m given he is naturally discharged. Similarly, let p_m^D and $1/\mu_m^{R,D}$ be the probability of readmission and expected readmission LOS of patient class m given he is demand-driven discharged. By Chen et al. (1998), we expect to have $p_m^N < p_m^D$ and $\mu_m^{R,N} > \mu_m^{R,D}$. Then the increase in readmission load attributable to the demand-driven discharge is precisely

$$\Delta\text{-readmission load} = \frac{p_m^{R,D}}{\mu_m^{R,D}} - \frac{p_m^{R,N}}{\mu_m^{R,N}}.$$

We will in the subsequent sections consider a priority rule that measures patient criticality via the Δ -readmission load score. In addition to fitting in with the ethos of a priority rule that can be interpreted as a criticality measure,

we see that this rule is consistent with assuming, in the notation of the previous sections, a one period cost function $C(s, A)$ that corresponds to the increase in readmission load because of the demand-driven discharge decision. In the appendix, we show that such a cost metric is also explicitly aligned with the desire to avoid a loss of throughput because of congestion effects.

Other measures of criticality. Although we have outlined the two broad criticality measures one might consider in the medical community, yet other measures have been proposed in the operations research community. In particular, Dobson et al. (2010) consider prioritizing patients based on a patients expected length of *remaining* stay. Unfortunately, this is a fairly difficult quantity to estimate and as such models to predict this quantity are also unavailable. For completeness, we will also consider this measure in our empirical investigation.

5. Empirical Data

The goal of this section is to calibrate a model from real data that will permit us to compare the clinically relevant policies discussed in the preceding section. We analyze patient data from seven different private hospitals for a total of 5,398 patients who completed at least one ICU visit.

Patient classes. Our first goal is to classify patients into a small number of groups, each of which is defined on the basis of physiological variables. There are many ways of doing this, and we chose a method that is aligned with the current process design philosophy of the hospital system from which the data for this study was obtained. In particular, we classified patients into five different classes based on the Kaiser Permanente Inpatient Risk Adjustment Score (see Escobar et al. 2008), which is a severity score used to predict the likelihood of death. These severity scores are based on a number of different factors including age; primary condition (cardiac, pneumonia, GI bleed, seizure, cancer, etc.); lab results obtained 72 hours prior to hospital admission; and chronic ailments (diabetes, kidney failure, etc.), and others. They are quite similar to the well-studied APACHE and SAPS scoring systems (for instance, the c statistic for this score is in the 0.88 range) with the important addition that they incorporate additional physiological information obtained for patients in this particular hospital system within a short time prior to their being admitted to the hospital (that APACHE or SAPS scores would not assume available). Like scoring rules of this type, the severity scores we use to classify patients may be interpreted as a mortality risk figure. This severity score is used in the hospital system we study, although the APACHE and SAPS scores are not available to us. We quantize these severity scores into one of five different bins of equal size. Table 1 summarizes the severity scores for the five patient classes as well as the percentage of survivors. It is important to note that we only use these scores as a convenient and clinically interpretable way of classifying patients. We do not

Table 1. Patient classes.

Patient class	Range for predicted mortality	No. of data points	% survivors
1	[0, 0.0048)	1,089	99.5
2	[0.0048, 0.0148)	1,084	97.0
3	[0.0148, 0.039)	1,097	94.7
4	[0.039, 0.1025)	1,067	91.8
5	[0.1025, 1)	1,061	85.4

use the severity score of a patient for the purposes of predicting mortality, length of stay, probability of readmission, and so-forth; rather, we directly estimate all of these factors from data.

ICU occupancy levels. Our data set indicates the utilization of the ICU upon patient discharge. We define the “near capacity” or “full” state as when the ICU occupancy level is at least 75% of its maximum. If the ICU occupancy is less than 75% of maximum, we say the ICU is in the “low” state. This characterization is similar to that in Kc and Terwiesch (2012) and acceptable from a medical perspective.

Sampling bias. Our study rests on the assumption that the statistics governing a patient’s length of stay in the ICU, the likelihood of their death, the likelihood of their readmission, and the lengths of any subsequent visits depend solely on their health condition as summarized by their Kaiser Permanente Inpatient Risk Adjustment Score and whether or not they were discharged from a full ICU. Because we are interested in isolating the impact of demand-driven discharge to accommodate new patients on patient length-of-stay statistics and the likelihood of readmission, it is important to check that the distribution of severity scores for patients in the group of patients discharged from a full ICU is close to that of patients discharged from an ICU in the low state. To this end, we use the Kolmogorov-Smirnov two-sample test (see Smirnov 1939 and related references), which is the continuous version of the chi-squared test. For each pair of ICU occupancy levels (full versus low), we compare the empirical distributions of severity using the Kolmogorov-Smirnov test to see if the samples come from the same distribution. We find that with significance level of 1%, the samples do come from the same distribution. Hence, we conclude with high probability, that the ICU occupancy level parameter and the severity scores of data points in our data set are independently distributed.

To summarize, a data point in our data set can be expressed as a tuple of the form $(S, D, (L_1, F_1), (L_2, F_2), \dots, (L_k, F_k))$, where S is a severity score, D is an indicator of patient death during hospital stay, L_n is the patient length of stay on his n th visit to the ICU in the episode, and F_n is an indicator for whether the ICU was full upon his n th discharge.

5.1. Estimation

We first estimate the probability of death for patients discharged from a low versus full ICU. We estimate the

nominal probability of death, $P(D | \text{Low})_m$, using the fraction of patients who were discharged from a low occupancy ICU and died during the same hospital stay:

$$P(D | \text{Low})_m = \frac{\sum_i \mathbf{1}_{\{D^i=1\}} \mathbf{1}_{\{F_1^i=0\}} \mathbf{1}_{\{S^i \in m\}}}{\sum_i \mathbf{1}_{\{F_1^i=0\}} \mathbf{1}_{\{S^i \in m\}}},$$

where $\{F_1^i = 0\}$ is the event that the ICU occupancy level was low upon discharge of patient i from his first ICU discharge and $\{S^i \in m\}$ is the event that the severity score of patient i defines him as class m . Similarly, we can calculate the probability of death when discharged from a full ICU:

$$P(D | \text{Full})_m = \frac{\sum_i \mathbf{1}_{\{D^i=1\}} \mathbf{1}_{\{F_1^i=1\}} \mathbf{1}_{\{S^i \in m\}}}{\sum_i \mathbf{1}_{\{F_1^i=1\}} \mathbf{1}_{\{S^i \in m\}}}.$$

Table 2 summarizes the estimated probabilities of death for each patient class along with the 95% confidence interval for these estimates.

We notice that it is difficult to discern any substantial impact of a demand-driven discharge on mortality. This is not particularly surprising: although there exist studies that suggest that demand-driven discharges increase mortality rates (for example Chrusch et al. 2009), there are others that find that mortality risks are not predicted by occupancy levels (Iwashyna et al. 2000).

Our estimator for the nominal length of stay (LOS) for patient type m , is simply the empirical average

$$\mu(\text{LOS}_{\text{low}}^0)_m = \frac{\sum_i L_1^i \mathbf{1}_{\{F_1^i=0\}} \mathbf{1}_{\{S^i \in m\}} \mathbf{1}_{\{D^i=0\}}}{\sum_i \mathbf{1}_{\{F_1^i=0\}} \mathbf{1}_{\{S^i \in m\}} \mathbf{1}_{\{D^i=0\}}}.$$

where $\{F_1^i = 0\}$ is the event that the ICU occupancy level was low upon discharge of patient i from his first ICU discharge and $\{S^i \in m\}$ is the event that the severity score of patient i defines him as class m . Similarly $\sigma(\text{LOS}_{\text{low}}^0)_m$ is an empirical standard deviation. Note that when calculating LOS, we exclude patients who died. This is common practice in the medical community because various factors, such as do-not-resuscitate orders can skew LOS estimates for patients who die (Norton et al. 2007, Rapoport et al. 1996).

We also calculate the fraction of these patients who return to the ICU during the same hospital stay to calculate a nominal probability of readmission, $P(R | \text{Low})_m$. These readmitted patients relapse because of numerous medical reasons unrelated to being discharged; the discharge is

likely to be a natural departure as there is no need to discharge patients in order to accommodate new ones when the ICU occupancy level is low and there are available beds. Thus,

$$P(R | \text{Low})_m = \frac{\sum_i \mathbf{1}_{\{L_2^i > 0\}} \mathbf{1}_{\{F_1^i=0\}} \mathbf{1}_{\{S^i \in m\}}}{\sum_i \mathbf{1}_{\{F_1^i=0\}} \mathbf{1}_{\{S^i \in m\}}},$$

where $\{L_2^i > 0\}$ denotes the event that patient i was readmitted.

Finally, of patients readmitted to the ICU from among those initially discharged from a nonfull ICU, we compute an estimate of their expected length of stay upon readmission, according to

$$\mu(\text{LOS}_{\text{low}}^R)_m = \frac{\sum_i L_2^i \mathbf{1}_{\{F_1^i=0\}} \mathbf{1}_{\{L_2^i > 0\}} \mathbf{1}_{\{S^i \in m\}} \mathbf{1}_{\{F_2^i=0\}} \mathbf{1}_{\{D^i=0\}}}{\sum_i \mathbf{1}_{\{F_1^i=0\}} \mathbf{1}_{\{L_2^i > 0\}} \mathbf{1}_{\{S^i \in m\}} \mathbf{1}_{\{F_2^i=0\}} \mathbf{1}_{\{D^i=0\}}},$$

where $\{F_2^i = 0\}$ denotes the event that patient i was discharged from a low occupancy ICU upon his second ICU discharge. Again, we exclude patients who died in this estimation. Notice that $\mu(\text{LOS}_{\text{low}}^R)_m$ is an estimate of patient length of stay upon readmission when the readmission is due to medical factors unrelated to demand-driven discharge. Table 3 states the values of the estimates for our data set including information about the relevant number of data points.

We compute similar estimates for patients discharged from a full ICU; we assume these discharges are demand driven. Of particular interest is the probability of patient readmission when a patient is discharged from a full ICU, $P(R | \text{Full})_m$. We estimate this probability according to

$$P(R | \text{Full})_m = \frac{\sum_i \mathbf{1}_{\{F_1^i=1\}} \mathbf{1}_{\{L_2^i > 0\}} \mathbf{1}_{\{S^i \in m\}}}{\sum_i \mathbf{1}_{\{F_1^i=1\}} \mathbf{1}_{\{S^i \in m\}}}.$$

We have seen that patients who are not discharged in order to accommodate new patients may require readmission (Table 3); we expect that patients who are discharged from a full ICU may require readmission for those same reasons *in addition* to complications that arise because of being demand-driven discharged. Therefore, we expect the probability of readmission when discharged from a full ICU to be higher than when discharged from a low ICU. We also estimate the expected length of stay of such readmitted patients according to

$$\mu(\text{LOS}_{\text{full}}^R)_m = \frac{\sum_i L_2^i \mathbf{1}_{\{F_1^i=1\}} \mathbf{1}_{\{L_2^i > 0\}} \mathbf{1}_{\{S^i \in m\}} \mathbf{1}_{\{F_2^i=0\}} \mathbf{1}_{\{D^i=0\}}}{\sum_i \mathbf{1}_{\{F_1^i=1\}} \mathbf{1}_{\{L_2^i > 0\}} \mathbf{1}_{\{S^i \in m\}} \mathbf{1}_{\{F_2^i=0\}} \mathbf{1}_{\{D^i=0\}}}.$$

Table 4 states the values of these estimates for our data set including information about the relevant number of data points.

Table 2. Mortality: Probability of death when patients naturally depart and when patients are demand-driven discharged.

Patient class	No. of data points	$P(D \text{Low})$	[95% CI]	No. of data points	$P(D \text{Full})$	[95% CI]
1	739	0.005	[0.000, 0.010]	350	0.003	[0.000, 0.009]
2	682	0.022	[0.011, 0.033]	402	0.017	[0.004, 0.030]
3	679	0.059	[0.041, 0.077]	418	0.043	[0.024, 0.062]
4	669	0.079	[0.059, 0.099]	398	0.088	[0.060, 0.116]
5	621	0.167	[0.138, 0.196]	440	0.116	[0.086, 0.146]

Table 3. Nominal patient parameters: Operational parameters when patients naturally depart and are not discharged in order to accommodate new patients.

Patient class	No. of data points	$\mu(\text{LOS}_{\text{low}}^0)$ (hours)	$\sigma(\text{LOS}_{\text{low}}^0)$	$P(R \text{Low})$	[95% CI]	No. of data points	$\mu(\text{LOS}_{\text{low}}^R)$ (hours)	$\sigma(\text{LOS}_{\text{low}}^R)$
1	735	45.7	134.2	0.073	[0.054, 0.092]	34	36.1	40.5
2	667	46.7	50.8	0.095	[0.073, 0.117]	46	66.0	118.1
3	639	59.7	98.4	0.102	[0.079, 0.125]	39	106.9	212.5
4	616	78.1	201.8	0.115	[0.091, 0.139]	45	110.5	289.3
5	517	89.6	116.7	0.119	[0.094, 0.115]	34	161.4	365.5

Notes. Average initial length of stay ($\text{LOS}_{\text{low}}^0$), readmission probability $P(R | \text{Low})$ and readmission length of stay ($\text{LOS}_{\text{low}}^R$) when discharged from a “low” occupancy ICU. Length of stay is given in hours.

Contrasting the results in Tables 3 and 4 we see that patients discharged during times of heavy ICU utilization are markedly more likely to be readmitted, all else being the same. In the following section, we will use the estimates we have computed here to construct and simulate the clinically relevant policies discussed in the previous section.

6. Performance Evaluation

The goal of this section is to explicitly construct the clinically relevant policies discussed in §4 using the estimates of the previous section. For each of the policies we construct, we will primarily be interested in characterizing two things.

Mortality. This is a first order measure of the clinical impact of any demand-driven discharge practice. Given our discussion in §4, one would hope that any of the clinically relevant discharge policies considered there results in effectively equivalent mortality rates. If this were not the case, this would be cause to question the clinical viability of the policies.

Measures of access. Assuming that two given policies possess similar mortality rates, one may be concerned about finer grained measures of performance. An important issue raised in §4—and indeed a focus of this paper and recent healthcare reform—was that of access. It is crucial that the demand-driven discharge policies employed ensure *equitable* and *maximal* access to ICU resources while of course, ensuring no sacrifice in terms of mortality rates. In fact, it is entirely within reason that these two goals are aligned as opposed to being at odds with each other.

Table 4. Demand-driven discharged patient parameters: Operational parameters when patients are discharged in order to accommodate new patients.

Patient type	No. of data points	$\mu(\text{LOS}_{\text{full}}^0)$ (hours)	$\sigma(\text{LOS}_{\text{full}}^0)$	$P(R \text{Full})$	[95% CI]	No. of data points	$\mu(\text{LOS}_{\text{full}}^R)$ (hours)	$\sigma(\text{LOS}_{\text{full}}^R)$
1	349	54.3	138.9	0.086	[0.057, 0.115]	9	61.4	71.8
2	395	51.7	54.1	0.109	[0.079, 0.140]	16	112.0	200.2
3	400	59.4	79.3	0.120	[0.089, 0.151]	17	99.6	86.8
4	363	62.8	68.1	0.136	[0.102, 0.170]	17	175.7	375.1
5	389	92.7	138.2	0.132	[0.100, 0.164]	17	237.1	577.8

Notes. Average initial length of stay ($\text{LOS}_{\text{full}}^0$), readmission probability $P(R | \text{Full})$, and readmission length of stay ($\text{LOS}_{\text{full}}^R$) when discharged from a “full” ICU. Length of stay is given in hours.

We next specify each of the policies discussed qualitatively in §4.

Mortality risk “P(D)”. Under this policy, if a demand-driven discharge is called for, one selects a patient from the class with the smallest probability of death, $P(D)$, of the patients currently in the ICU. Table 2 calibrates these figures for patients in our data set. This translates to the order 1, 2, 3, 4, 5.

Readmission risk I “P(R)”. Under this policy, one selects a patient from the class with the smallest nominal probability of readmission, $P(R)$, of the patients currently in the ICU. In particular, given the estimates from our data set reported in Table 3, this translates to the order 1, 2, 3, 4, 5.

Readmission risk II “Δ-load”. This policy, which as discussed earlier, is a refinement of the readmission risk metric above, has been a focal point of our study. We can estimate the increase in readmission load for a given patient class, m , as the quantity

$$P(R | \text{Full})_m \mu(\text{LOS}_{\text{full}}^R)_m - P(R | \text{Low})_m \mu(\text{LOS}_{\text{low}}^R)_m.$$

Using the data from Tables 3 and 4, this translates to the priority order 3, 1, 2, 4, 5.

Remaining length of stay “LOS”. Under this policy, one selects a patient from that class with the smallest *remaining* length of stay. As such this is not a static index rule. In particular, one needs to compute, for a patient of class m that has been in the ICU for time t , the quantity $\mathbb{E}[\text{LOS}_{\text{low}}^0 | \text{LOS}_{\text{low}}^0 \geq t]$, and prioritize patients in increasing order of this quantity. In our simulations, we give this policy more power and assume the realization for ICU LOS is known as

Table 5. Estimated parameters for discharge policies.

Patient type	Nominal $P(D)$	Nominal $P(R)$	Δ -readmission load (hours)
1	0.005	0.073	2.65
2	0.022	0.095	5.94
3	0.059	0.102	1.05
4	0.079	0.115	11.19
5	0.167	0.119	12.09

soon as a patient begins ICU care. This policy is analyzed in Dobson et al. (2010) albeit for a model that is agnostic to readmission loads.

Table 5 summarizes the patient parameters for the first three policies. It is interesting to note that of the first three policies, all three policies choose to protect patients of types 4 and 5 from a demand-driven discharge. These are patients with relatively higher mortality risk, and as such this is a desirable feature. Interestingly, the Δ -load policy differs from the first two in how it prioritizes the first three patient classes that have low mortality risk. This allows for the following interpretation of the Δ -load policy: it ensures that patients with high mortality risk are the least likely to be subject to a demand-driven discharge while carefully prioritizing among patients with low mortality risk to account not only for the likelihood they would have to be readmitted as a consequence of the discharge, but also the extent of the care they might require if such a readmission were to occur.

We consider the following simulation setup: We assume a time horizon of one week where admission and discharge decisions are made every six minutes, or 10 times within an hour, and consider an ICU with $B = 10$ beds. Although these decisions may in reality occur on a continuous basis, patient transfers are not instantaneous and the granularity of six minutes per hour is fine enough to emulate an actual ICU. Discharge policy simulations are over 1,000 sample paths each. We use the parameters estimated in Table 5 for nominal length of stay, probability of death, probability of readmission, and change in expected readmission load. A patient’s nominal length of stay is log-normally distributed. We vary the probability of an arrival, λ between 0 and 0.021 (i.e., between zero and five arrivals on average every 24 hours). An arrival rate $\lambda = 0.021$ corresponds to five patient per day, i.e., a turnover of 1/2 the beds in the ICU each day, which is about the highest load seen in the ICU. We use a uniform traffic mix across patient classes, which is consistent with the empirical data. We note that in this numerical study, we do not include diversions and disease progression as patient data required to develop these models is fairly limited. For instance, Gajic et al. (2008) is one of the few existing works that try to predict readmission risks, and it is a static model. We leave such studies for future research as more patient data becomes available. We now report on the two issues we set out to examine, namely, mortality and patient access.

6.1. Mortality Rates

We compare the number of deaths per week under the various discharge policies. We consider an arrival rate of 2.5 patients per day, which corresponds to the load an average hospital could expect to a 10 bed ICU. In column (a) of Table 6 we compare the number of deaths per week using the point estimates of $P(D | \text{full})$ and $P(D | \text{low})$ given in Table 2. We also consider the following robustness check using the confidence intervals computed for our class specific mortality rate estimates: we consider that the various probabilities (namely, $P(D | \text{full})_m$ and $P(D | \text{low})_m$) each take on one of their upper or lower confidence limits, and consider all the 2^{10} resulting parameter combinations. We conduct a separate simulation for each of these parameter combinations and report for each discharge policy the lowest and highest mortality rates across parameter combinations. The results are summarized in Table 6. We can see that using both the point estimates, as well as under our robustness check, all three policies are remarkably similar.

We next consider a further robustness check assuming that, in fact, the probability of death upon being demand-driven discharged is substantially increased (beyond the value estimated in the data)—we set the probability of death for a demand-driven discharged patient 10%, 20%, 30%, 40%, and 50% higher than the estimated probability of death for that patient class given in Table 2. We compare the relative increase (decrease) in the number of deaths compared to the proposed Δ -readmission load policy. Table 7 summarizes these results. Again, the table reveals that the three policies continue to remain essentially identical across the range of perturbations with no single policy dominating.

From these experiments, we conclude that in as much as mortality rates are concerned all four policies are viable and result in essentially identical mortality rates. In spite of the fact that the policies differ from each other, this reaffirms our earlier assertion that all four of the policies will protect patients with high mortality rates from a demand-driven discharge.

Table 6. Weekly mortality rate using (a) point estimates, (b) the combination over the 95% confidence intervals with the lowest rate, and (c) the combination over the 95% confidence intervals with the highest rate.

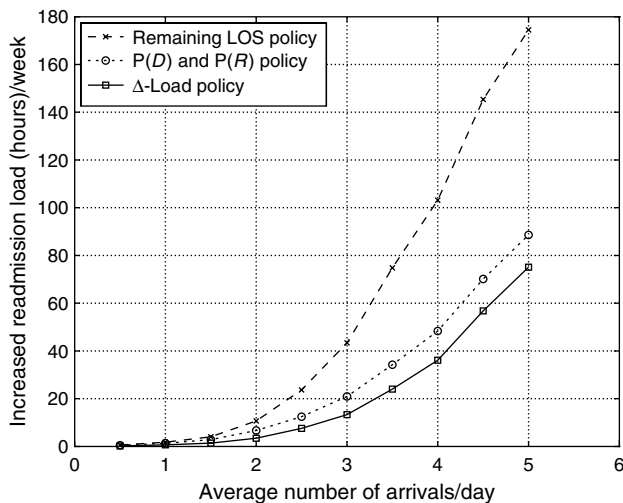
Policy	(a) No. of deaths	(b) Min. no. of deaths	(c) Max. no. of deaths
Δ -readmission load	1.014	0.751	1.325
$P(\text{death})$ and $P(\text{readmission})$	1.004	0.764	1.332
Shortest remaining LOS	1.022	0.740	1.303

Table 7. Percentage increase over Δ -readmission load policy in weekly mortality rate when artificially inflating $P(\text{death} | \text{full})$.

Inflation factor (%)	Shortest remaining LOS (%)	$P(\text{death})$ and $P(\text{readmission})$ (%)
0	-0.9	0.9
10	0.0	0.5
20	0.8	0.1
30	1.5	-0.4
40	2.2	-0.7
50	2.6	-1.2

6.2. Patient Access

We measure access via the following proxy: Because demand-driven discharges result in an increase in the expected critical care requirements for the discharged patient down the road, we measure the expected increase in these requirements, measured in hours of ICU care. In particular, we measure the expected increase in readmission load incurred because of demand-driven discharges under all four policies. Figure 2 shows the expected increased readmission load in hours for the four discharge policies. We can see that the proposed Δ -load policy outperforms each of the benchmarks—in some cases by nearly 30%. The next best policy in this regard is the one based on (unadjusted) readmission and mortality risks, i.e., the $P(R)$ and $P(D)$ index policy. Thus, although the problem of minimizing readmission load because of required demand-driven discharges is a hard one, the proposed Δ -load policy appears to substantially outperform the benchmarks studied here. As the arrival rate increases, more patients will need

Figure 2. Performance of proposed index policy compared to benchmarks for various arrival rates and distribution across patient types according to the proportions seen in the empirical data.

to be demand-driven discharged in order to accommodate the high influx of new patients. Consequently, the expected readmission load increases significantly.

To appreciate the physical meaning of the costs estimated in these experiments, we note that with 24 hours in a day, an additional cost of $24 \times 7 = 168$ hours corresponds to the loss of an entire bed for one week because it will be occupied by readmitted patients. What we see is that for five patient arrivals per day, the Δ -load policy incurs readmission load that is 13.5 hours lower than the next best policy (the $P(D)$ and $P(R)$ policy), which corresponds to the loss of a single ICU bed (in a 10 bed ICU) for a little more than half a day per week. Over the course of a year this corresponds to a free ICU bed for nearly a whole month. The savings relative to the LOS index policy are higher. Finally, in light of our study on mortality rates, these differences in performance do not come at the cost of increased mortality.

In conclusion, we have observed the following:

Mortality. All four policies we considered at the outset in §4 result in essentially identical mortality rates. We have verified this fact across multiple “robustness” checks. We attribute this to an attractive feature common to the first three policies, namely, the fact that patients with high mortality risk are protected from a demand-driven discharge.

Access. In terms of access (or equivalently, increase in ICU load incurred because of demand-driven discharges) the policies are quite dissimilar. The Δ -load policy (that has been a focal point in this paper) provides the greatest access. We attribute this to the fact that the policy carefully prioritizes among patients with low mortality risk.

As such, we believe that the Δ -load policy might serve as a useful guide to intensivists prioritizing demand-driven discharge decisions among patients medically fit for discharge.

7. Conclusion

Faced with the need to accommodate an acute, newly admitted patient, a clinician may select from among patients currently in the ICU, a relatively “stable” patient for transfer to a less richly staffed hospital unit. A patient so discharged from the ICU faces risks of physiological deterioration that may ultimately require readmission to the ICU. This is, of course, not an ideal situation either from an efficiency standpoint or the standpoint of ideal patient outcomes. The present work studied the *feasibility* of developing a decision support tool to aid clinicians in these difficult decisions. We have attempted to gauge the *value* of such a support tool using a large patient flow data set and quantified this value in terms of potential reductions in readmitted patient load.

The model we have developed revolves around simple estimates of the cost associated with a demand-driven patient discharge. We examine a number of clinically relevant cost metrics including mortality and readmission risks. We focus on a measure of readmission risk that

incorporates the likelihood of readmission in addition to the complexity of the readmission: change in readmission load. We estimated our model from actual patient-flow data. Given our model, we developed a simple index based policy to serve as a decision support tool to a physician making the aforementioned discharge decisions. Our support tool is, by its structure, easy to implement from a clinical standpoint, and highly robust to estimation errors. The latter point is well reflected in our empirical study. Our study suggests that implementation of our support tool could result in substantial reductions in readmitted patient load without sacrificing mortality even under modest assumptions on patient traffic, at least in the context of the hospital system from which we collected the data for the study. Although the Kaiser Permanente Inpatient Risk Adjustment Score has been shown to have similar predictive power as the widely used SAPS and APACHE scores (Escobar et al. 2008), the value of our approach has not been established for these other severity scores. This represents an interesting avenue for future empirical investigation. It is remarkable that our model demonstrates benefits despite (from a clinical standpoint) being relatively simple—for example, it does not include diagnostic or physiologic data available at the time that a patient was discharged the following:

This work suggests several future potential research directions, including the following:

1. Developing more complex predictive models of patient dynamics that recognize the evolution of patients over the course of their stay. We believe that the present study is sufficient motivation to collect data that would allow us to identify such a model. Such data could be employed to assign patients a “readiness for discharge” severity score similar in concept to other existing severity of illness scores. This is also key to practical deployment of a decision support tool.

2. It would be interesting to understand the impact of a demand-driven discharge on other quantities of interest, particularly metrics measuring quality of life impact.

3. Theoretically, we have shown that our index policy is optimal in certain regimes and guaranteed to incur readmission loads of no greater than a factor of $(\hat{\rho} + 1)$ of an optimal policy in general. It would be interesting to understand traffic regimes where this bound could be made tighter—this is, of course, a somewhat secondary pursuit but nonetheless very interesting from a theoretical perspective.

4. It would be interesting to initiate a study of ICU admissions so as to move toward a more holistic view of equitable and optimal allocation of hospital resources.

Appendix A. Greedy Suboptimality

Consider the case with $B = 2$ beds and a time horizon of $T = 2$. There are two patient types, $i \in \{1, 2\}$. The parameters for each patient type are as follows for some small $\epsilon > 0$:

$$\text{for } i = 1: \mu_1^0 = 1/2, \quad \phi_1 = 1$$

$$\text{for } i = 2: \mu_2^0 = 1, \quad \phi_2 = 1 - \epsilon.$$

Therefore, patient type 1 has nominal expected length of stay of 2 and cost of 1. Similarly, patient type 2 has nominal expected length of stay of 1 and cost of $1 - \epsilon$.

Consider an initial state at $t = 0$ such that there exists two ICU patients: one of each type. Hence, $x_{0,1} = 1$ and $x_{0,2} = 1$. Also, a new patient of type 1 arrives at $t = 0$ and $t = 1$, i.e., $y_{0,1} = y_{1,1} = 1$ and $y_{0,2} = y_{1,2} = 0$.

At $t = 0$, there are already two patients in the ICU, and a new patient arrives. Therefore, a current patient must be discharged in order to accommodate the new patient. The greedy policy discharges patient type 2 at $t = 0$ because its cost is less than that of patient type 1. This comes at a cost of $1 - \epsilon$. Now, with this demand-driven discharge and the admission of the new patient there are two type 1 patients occupying the ICU. With probability $1/4$ neither type 1 patient completes service and departs by $t = 1$ and with the second new arrival, a patient must be discharged to accommodate this new arrival at a cost of 1. With probability $3/4$ at least one of the type 1 patients completes service prior to the second new arrival and no demand-driven discharge is required at $t = 1$. Hence, the expected cost of the greedy policy is $1 - \epsilon + 1/4 = 5/4 - \epsilon$.

On the other hand, the optimal policy recognizes that patient type 2 has a very short length of stay and decides not to discharge this patient at $t = 0$. Instead the optimal policy discharges patient type 1 to accommodate the new patient, incurring a cost of 1. Now with this demand-driven discharge and the admission of the new patient, there is one type 1 patient and one type 2 patient occupying the ICU. At the end of time slot $t = 0$, the type 2 patient completes service and departs naturally with probability 1. Regardless of whether the type 1 patient departs naturally, when the second new arrival comes at $t = 1$, it can immediately be accommodated without requiring a demand-driven discharge of a current patient. Hence, the expected cost of the optimal policy is 1.

Taking $\epsilon \rightarrow 0$ we see that $J^*(s_0) \leq (4/5)J^s(s_0)$ here.

Appendix B. A Connection with Throughput

Here we make precise the connection with throughput maximization when the cost metric of interest is the Δ -readmission load associated with a demand-driven discharge. We consider a stylized model of the ICU that accounts for patient readmissions. Patients who are naturally discharged require ICU readmission with probability 0. Patients who are demand-driven discharged are readmitted to the ICU with probability p_m and have readmission LOS that is exponentially distributed with mean $1/\mu_m^R$. Hence, the cost associated with a demand-driven discharge of patient type m , is the Δ -readmission load:

$$C(s, m) = \frac{p_m}{\mu_m^R}.$$

Consider an ICU with C beds. We consider the following setup:

1. There are B beds reserved for first-time arrivals with $C - B \triangleq B'$ beds reserved for readmissions. Any reference to “state” will be understood to correspond to the occupants of these B beds and we will consequently employ the notation in §2.

2. The readmission queue is served according to a first-in-first-out discipline.

3. In the event that B beds are occupied by first-time visitors, a new arrival will prompt a demand-driven discharge according to a stationary policy π that monitors the state of the B beds reserved for first-time arrivals.

Note that readmitted patients cannot be demand-driven discharged. The rationale for this is natural: readmitted patients are typically much worse off and have higher mortality rates and longer lengths of stay. This is well established in the medical literature (see Chen et al. 1998, Durbin and Kopel 1993, Snow et al. 1985 among others). As such, subjecting such patients to a demand-driven discharge is likely to be highly undesirable from a practitioners perspective. In addition, the policy that prioritizes patients should a demand-driven discharge be required may only consider the state of the B beds reserved for first time arrivals; one may dispense with this restriction but doing so is beyond our scope here.

Given a vector $\lambda \in [0, 1]^M$ defined so that $\lambda a_{t,m} \triangleq \lambda_m$ for all m (assuming time homogenous rates), we will refer to a policy π as *stabilizing* for λ if under this policy the readmission queue is stable. More precisely, we require the sequence of waiting times $\{W_n\}$ experienced by patients in the readmission queue (a waiting time is defined in the usual sense as the time between entry into the readmission queue and the time before service begins), has a sub-sequence that converges to a random variable W that is a.e. finite.

Now, let us denote by the sequence T_n the interarrival time between the n th and $(n+1)$ st entry to the readmission queue, and by S_n , the service time required by the n th patient. Assume moreover that no demand-driven discharges occur in the absence of a need for one, i.e., $\pi(s) = 0$ if $s \notin \{(x, y) : \sum_m x(s)_m + y(s)_m = B + 1, \sum_m y(s)_m = 1\} \triangleq \hat{\mathcal{S}}_{\text{full}}$ (recall again, that s here corresponds to the state of the B beds reserved for first-time admissions). Then, T_n is simply the time between the n th and $(n+1)$ st visit to a state in the set $\hat{\mathcal{S}}_{\text{full}}$ while S_n is a geometric ($\mu_{\pi(s_n)}^R$) random variable with probability $p_{\pi(s_n)}$ (where s_n corresponds to the state of the B beds for first-time arrivals upon the n th discharge) and zero with the remaining probability. Now, if $s_0 \sim \nu_\pi$, then it is not hard to see that $\{T_n, S_n\}$ is a stationary process. The process is also ergodic; a consequence of the ergodicity of the Markov chain induced by π . A classical result of Loynes (Theorem 8 of Loynes 1963) then establishes that the readmission queue is stable if $E[T_0] > E[S_0]/(C - B)$, and unstable if $E[T_0] < E[S_0]/(C - B)$. Now, elementary arguments (see Durrett 1996) can be used to show that $E[T_0] = 1/\sum_{s \in \hat{\mathcal{S}}_{\text{full}}} \nu_\pi(s)$ and $E[S_0] = \sum_{s \in \hat{\mathcal{S}}_{\text{full}}} \nu_\pi(s)C(s, \pi(s))/\sum_{s \in \hat{\mathcal{S}}_{\text{full}}} \nu_\pi(s)$. In other words, we have that the readmission queue is stable if

$$\kappa^\pi < C - B,$$

and unstable if $\kappa^\pi > C - B$, so that minimizing κ^π maximizes throughput, which motivates the problem that is the focus of our study.

In addition, the following theorem shows that heuristics for the problem of minimizing long-run readmission costs incur a proportionate “dilation” of the set of arrival rate profiles that will result in stable readmission queues. In particular, let λ be a vector of arrival rates that is in the interior of the throughput region for our model. By this we understand that there exists a demand-driven discharge policy π_λ^* under which the readmission queue is stable when the arrival rate vector is λ , and moreover there exists an $\epsilon > 0$ such that the arrival rate vector $\lambda(1 + \epsilon)$ can also be stabilized. Let us denote by $\pi_{\alpha\lambda}^*$ a policy minimizing κ^π for the arrival rate vector $\alpha\lambda$, where $\alpha \in (0, 1]$. Finally, let $\hat{\pi}_{\alpha\lambda}$ be a possibly suboptimal demand-driven discharge policy for the arrival rate $\alpha\lambda$ satisfying $\kappa^{\hat{\pi}_{\alpha\lambda}}/\kappa^{\pi_{\alpha\lambda}^*} \leq 1/\alpha$. We have the following:

THEOREM 5. *Assuming an arrival rate vector $\alpha\lambda$, the readmission queue is stable under the demand-driven discharge policy $\hat{\pi}_{\alpha\lambda}$.*

PROOF. Let us denote by $\pi_{\alpha\lambda}^*$ (respectively, π_λ^*) a policy minimizing κ^π in a system with arrival rate vector $\alpha\lambda$ (respectively, λ). Now consider the following suboptimal policy for an arrival rate $\alpha\lambda$: we simulate arrivals of “fictitious” patients, so that the net stream of patients (both actual and fictitious) has arrival rate λ . To this system we apply policy π_λ^* . Now by construction, a discharge under this policy will correspond to the discharge of an actual patient with probability α ; with the remaining probability, the discharge will be one of a fictitious patient and incur no costs. It thus follows that this suboptimal policy incurs a cost of precisely $\alpha\kappa^{\pi_\lambda^*}$. Moreover, because it is suboptimal for the arrival rate vector $\alpha\lambda$, it must be that

$$\kappa^{\pi_{\alpha\lambda}^*} \leq \alpha\kappa^{\pi_\lambda^*}.$$

It follows that

$$\kappa^{\hat{\pi}_{\alpha\lambda}} \leq (1/\alpha)\kappa^{\pi_{\alpha\lambda}^*} \leq \kappa^{\pi_\lambda^*}.$$

But given the fact that λ was in the interior of the stability region, it must be (by our earlier argument that showed $\kappa^{\pi_{\alpha\lambda}^*} \leq \alpha\kappa^{\pi_\lambda^*}$) that $\kappa^{\pi_\lambda^*} < C - B^*$, so that $\kappa^{\hat{\pi}_{\alpha\lambda}} < C - B^*$, from which the claim follows. \square

We have demonstrated a stationary policy π^s satisfying, for a given arrival rate vector λ , $\kappa^{\pi^s}/\kappa^{\pi_\lambda^*} \leq 1/(1 + \hat{\rho})$, where $\hat{\rho}$ was a measure of utilization. It follows that should the readmission queue be *unstable* under π^s , then it will remain unstable for any arrival rate vector that strictly dominates $(1 + \hat{\rho})\lambda$ under *any* stationary discharge policy. In other words, the use of the π^s policy will correspond to a dilation of the throughput region by a factor corresponding to the approximation guarantee we have established.

Appendix C. Miscellaneous Technical Proofs

PROOF OF THEOREM 1. We will, without loss, consider states s at which all feasible actions require the demand-driven discharge of a current patient (who has not yet completed treatment); i.e., $\sum_m x(s)_m = B$ and $y(s) \neq 0$. For the sake of a contradiction, we will assume that under any optimal policy π^* , $\pi^*(s) \notin \arg \min_{m: x(s)_m > 0} \phi_m$, i.e., the patient selected for the demand-driven discharge under any optimal policy is not among the set of patient types that minimizes one-period costs at state s . For notational convenience, we take $\pi^*(s) = j$, and $i = \pi^s(s) \in \arg \min_{m: x(s)_m > 0} \phi_m$. Thus, by assumption we have that

$$\begin{aligned} J^*(s) &= C(s, j) + E[J^*(S(s, j))] \\ &< C(s, i) + E[J^*(S(s, i))]. \end{aligned} \quad (C1)$$

Now, let $s_i = S(s, j)$, and $s_j = S(s, i)$. We may assume that $x(s_i)_k = x(s_j)_k \forall k \neq i, j$. Moreover, because $C(s, i) < C(s, j)$, we have $1/\mu_i^0 \geq 1/\mu_j^0$ so that we may couple sample paths in the system that used the optimal policy in state s (demand-driven discharged patient j) with the system that used the greedy policy at state s (demand-driven discharged patient i) so that patient i finishes service and departs in the epoch subsequent to $t(s)$ in the former system only if j finishes service and departs naturally in that same epoch in the latter system. Thus, in time slot $t(s) + 1$ we have either that (i) $x(s_i)_i - x(s_j)_i = 1$ and $x(s_j)_j - x(s_i)_j = 0$, (ii) $x(s_i)_i - x(s_j)_i = 0$ and $x(s_j)_j - x(s_i)_j = 0$,

or (iii) $x(s_i)_i - x(s_j)_i = 1$ and $x(s_j)_j - x(s_i)_j = 1$. In case (i), Lemma 1 implies that $J^*(s_i) \geq J^*(s_j)$. In case (ii), we clearly have $J^*(s_i) = J^*(s_j)$ because $s_i = s_j$.

Let us consider case (iii), which says that neither patient i nor j have departed by time slot $t(s) + 1$. We couple the systems starting at states s_i and s_j so that they see identical arrivals and identical service times (departures) for the patients they have in common. Moreover, we couple the service times of the additional type i patient in the s_i system and the additional type j patient in the s_j system as follows: If after any required demand-driven discharges in a particular time step, patient i and j both remain in their respective systems, patient j will complete/depart with probability μ_j^0 . If patient j departs, patient i will depart in the same time step with probability μ_i^0/μ_j^0 ; if patient j does not complete, then neither will patient i . If only one of i or j are present, they will complete with probability μ_i^0 and μ_j^0 , respectively.

Now let us consider using the following suboptimal policy for the system starting at state s_j : we assume that the additional type j patient is in fact a type i patient, and apply the optimal policy for this transformed state. If at some point the type j patient completes service naturally, we choose to register this departure with probability μ_i^0/μ_j^0 , and with the remaining probability assume a “virtual” additional type i patient that will complete service in subsequent periods with probability μ_i^0 . If at some point the discharge policy chooses the additional type j patient (which it regards as a type i patient) for the demand-driven discharge, we charge ourselves $C(s, j)$ (notice that this may occur after the actual patient has already departed and correspond to the demand-driven discharge of the virtual patient), so that the costs incurred here are certainly higher than under an optimal policy for the s_j system. Call this policy π' . We use the optimal policy in the s_i system.

Let \bar{p}_i be the probability that the additional type i patient will have to be demand-driven discharged in the s_i system. Now we have that $J^*(s_i) = \bar{C} + \bar{p}_i C(s, i)$, where \bar{C} is the total readmission costs incurred for patients excluding the additional type i patient. Notice that under our coupling, $J^{\pi'}(s_i) = \bar{C} + \bar{p}_i C(s, j) = J^*(s_i) + \bar{p}_i [C(s, j) - C(s, i)]$. Consequently, we have that $J^*(s_j) - J^*(s_i) \leq \bar{p}_i [C(s, j) - C(s, i)]$.

Cases (i)–(iii) together yield $E[J^*(S(s, i)) - J^*(S(s, j))] \leq C(s, j) - C(s, i)$, which contradicts (C1) (because $C(s, i) \neq C(s, j)$) and yields our result. \square

PROOF OF LEMMA 1. Consider a coupling of the systems starting at state s and s' wherein both systems witness identical sample paths for patient arrivals and identical service times for the patients they have in common. More precisely, assuming that at time t , the systems are in states s_t and s'_t , respectively, the patients who arrive in both systems are coupled so that $y(s_t) = y(s'_t)$. Let $z(s_t)$ and $z(s'_t)$, be the patient vectors in both systems after these arrivals and any potential demand-driven discharges. Then the number of service completions in both systems over the remainder of the t th epoch are coupled as follows: If $z(s_t)_m \geq z(s'_t)_m$, then the number of patients of type m that finish service and depart naturally from the s' system is given by the binomial- $(z(s'_t)_m, \mu_m^0)$ random variable $X'_{t,m}$, and the number of patients of type m that finish service and depart naturally from the s system is given by $X'_{t,m} + Z_{t,m}$, where $Z_{t,m}$ is a binomial- $(z(s_t)_m - z(s'_t)_m, \mu_m^0)$ random variable. A symmetric situation must hold if $z(s'_t)_m \geq z(s_t)_m$.

Now assume that the system starting at s uses an optimal policy, whereas the system starting at state s' “mimics” the actions of

the s system (call this policy $\bar{\pi}$), so that if the s system chooses to demand-driven discharge a patient of a particular type, the s' system will also choose to discharge a patient of that type should such a patient be available, whether or not this demand-driven discharge is called for (i.e., whether or not a new patient has arrived and there are no available beds). In the event that the s' system needs to make a demand-driven patient discharge and the s system either does not need to make a demand-driven discharge or else selects to demand-driven discharge a patient of a class not available in the s' system, the s' system discharges a randomly chosen patient from among those available. It is easy to see that $\bar{\pi}$ is an admissible randomized nonanticipatory policy: starting at state s' one adds “virtual” patients so that the total number of patients (real and virtual) of a given type in the s' system are identical to the number in the s system. One then employs an optimal policy, and simulates service completion for virtual patients. We now show that under our coupling, $x(s_t) \geq x(s'_t)$ for all t .

The proof is based on induction in time. The base case follows from our assumption that $x(s) \geq x(s')$. We assume that for all $t \leq k$, $x(s_t) \geq x(s'_t)$ and will show this implies the same is true for $t = k + 1$. Let $A_k = \pi^*(s_k)$ and A'_k be the patient discharged at time k under the $\bar{\pi}$ policy. Note that $A'_{k,m} \leq A_{k,m}$ by our definition of $\bar{\pi}$ and the induction hypothesis. We have

$$\begin{aligned} x(s_{k+1})_m - x(s'_{k+1})_m &= [(x(s_k)_m + y(s_k)_m - A_{k,m})^+ - X_{k,m}] \\ &\quad - [(x(s'_k)_m + y(s'_k)_m - A'_{k,m})^+ - X'_{k,m}] \\ &\geq x(s_k)_m - x(s'_k)_m + X'_{k,m} - X_{k,m} \\ &= x(s_k)_m - x(s'_k)_m - Z_{k,m} \\ &\geq 0. \end{aligned}$$

The first inequality comes from our coupling and the definition of the two policies. The second inequality follows from the definition of $Z_{t,m}$; $Z_{t,m} \leq x(s_t)_m - x(s'_t)_m$.

We may thus establish that for all $t(s) \leq t \leq T$, $A_t \geq A'_t$, so that $C(s_t, \pi^*(s_t)) \geq C(s'_t, \bar{\pi}(s'_t))$ for all such t . Taking expectations over the random patient arrivals and departures, we have $J^*(s) \geq J^{\bar{\pi}}(s') \geq J^*(s')$, which is the result. \square

PROOF OF LEMMA 2. Without loss, we assume $\pi^*(s) \neq \pi^g(s)$ (else, there is nothing to prove). By definition, we must have $x(s)^{\pi^g(s)}, x(s)^{\pi^*(s)} > 0$. Let $\tilde{S}(s, \pi^*(s))$ be the next state obtained if one discharged both $\pi^*(s)$ and $\pi^g(s)$ in state s . In particular, we define $\tilde{S}(s, \pi^*(s)) \triangleq \tilde{s}$ according to

$$\begin{aligned} x(\tilde{s})^{\pi^*(s)} &= x(s)^{\pi^*(s)} + y(s)^{\pi^*(s)} - 1 - X_{t(s), \pi^*(s)}, \\ x(\tilde{s})^{\pi^g(s)} &= x(s)^{\pi^g(s)} + y(s)^{\pi^g(s)} - 1 - X_{t(s), \pi^g(s)}, \\ x(\tilde{s})_m &= x(s)_m + y(s)_m - X_{t(s), m}, \quad m \neq \pi^*(s), \pi^g(s) \\ y(\tilde{s})_m &= Y_{t(s)+1, m}, \\ t(\tilde{s}) &= t(s) + 1, \end{aligned}$$

where analogous to our earlier description of $S(s, a)$, we define $X_{t(s), \pi^*(s)}$ (respectively, $X_{t(s), \pi^g(s)}$) as a binomial $(x(s)^{\pi^*(s)} + y(s)^{\pi^*(s)} - 1, \mu_{\pi^*(s)}^0)$ (respectively, binomial $(x(s)^{\pi^g(s)} + y(s)^{\pi^g(s)} - 1, \mu_{\pi^g(s)}^0)$) random variable. For $m \neq \pi^*(s), \pi^g(s)$, we define $X_{t(s), m}$ as a binomial $(x(s)_m + y(s)_m, \mu_m^0)$ random variable; $Y_{t(s)+1, m}$ is defined as before for all m . Now, by construction, $x(\tilde{S}(s, \pi^*(s))) \leq x(S(s, \pi^*(s)))$, while

$y(\tilde{S}(s, \pi^*(s))) = y(S(s, \pi^*(s)))$, so that by Lemma 1, we have that $E[J^*(\tilde{S}(s, \pi^*(s)))] \leq E[J^*(S(s, \pi^*(s)))]$.

Now, let us consider the following suboptimal policy π' for the system in which the greedy action is taken at state s . Define $\tau = \min\{T > t > t(s) : \sum_m Y_{t,m} = 1\}$; i.e., τ is the first time after the current time step $t(s)$ that an arrival occurs (or infinite if no arrival occurs prior to time T). Then on the event that $x(s_{\tau})_{\pi^*(s)} = x(s)_{\pi^*(s)} + y_{t(s), \pi^*(s)} - 1$, π' simply takes the optimal action for $t \geq \tau$ (so that, in fact π' coincides with π^* on this event). On the event that $x(s_{\tau})_{\pi^*(s)} = x(s)_{\pi^*(s)} + y_{t(s), \pi^*(s)}$, $\pi'(s_{\tau}) = \pi^*(s)$, and π' takes actions according to the optimal policy π^* for $t > \tau$. The probability that an eviction occurs under π' at τ is simply the probability that no patient of type $\pi^*(s)$ has departed prior to the next arrival; an event whose probability is at most $\lambda/(\lambda + \mu_{\pi^*(s)}^0)$. Observe moreover that we may couple the systems starting at state $S(s, \pi^*(s))$ and $\tilde{S}(s, \pi^*(s))$ so that under the π' policy in the former system and the optimal policy in the latter, both state processes agree on $t > \tau$, and moreover, no eviction will be required at times $t \leq \tau$ in the latter system. It follows that

$$E[J^{\pi'}(S(s, \pi^*(s)))] \leq \frac{\lambda}{\lambda + \mu_{\pi^*(s)}^0} C(s, \pi^*(s)) + E[J^*(\tilde{S}(s, \pi^*(s)))]$$

Because $E[J^*(S(s, \pi^*(s)))] \geq E[J^{\pi'}(S(s, \pi^*(s)))]$ and as established earlier, $E[J^*(\tilde{S}(s, \pi^*(s)))] \leq E[J^*(S(s, \pi^*(s)))]$, the result follows. \square

Endnotes

1. Because a strict (one-to-two in California) nurse-to-patient ratio must be maintained, it is often the size of the nursing staff that determines the number of available ICU beds rather than the actual number of physical beds that are available.
2. We later consider an extension of our model that includes the additional option of blocking new patients.

Acknowledgments

This project was approved by the Kaiser Permanente Northern California Institutional Board for the Protection of Human Subjects. The authors wish to thank Marla Gardner, John D. Greene, and Juan Carlos LaGuardia for assistance in preparing study data sets, and Dr. Philip Madvig and Edward Thomas for administrative support and assistance. This work was partially supported by the Permanente Medical Group, Inc., Kaiser Foundation Hospitals, Inc., and the Sidney Garfield Memorial Fund [Grant 115-9518, "Early detection of impending physiologic deterioration in hospitalized patients"]. The second author was supported, in part, by a National Science Foundation CAREER Award [CMMI-1054034].

References

- Allon G, Deo S, Lin W (2012) Impact of size and occupancy of hospital on the extent of ambulance diversion: Theory and evidence. Working paper, Kellogg Graduate School of Management, Northwestern University, Evanston, IL.
- American Thoracic Society (1997) Fair allocation of intensive care unit resources. *Amer. J. Respiratory Critical Care Medicine* 156(4): 1282–1301.

- Bone RC, McElwee NE, Eubanks DH, Gluck EH (1993) Analysis of indications for intensive care unit admission. Clinical efficacy assessment project: American College of Physicians. *CHEST* 104:1806–1811.
- Chalfin DB (2005) Length of intensive care unit stay and patient outcome: The long and short of it all. *Critical Care Medicine* 33:2119–2120.
- Chalfin DB, Trzeciak S, Likourezos A, Baumann BM, Dellinger RP (2007) Impact of delayed transfer of critically ill patients from the emergency department to the intensive care unit. *Critical Care Medicine* 35:1477–1483.
- Chan CW, Farias VF (2009) Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems. *Math. Oper. Res.* 34(2):333–350.
- Chen LM, Martin CM, Keenan SP, Sibbald WJ (1998) Patients readmitted to the intensive care unit during the same hospitalization: Clinical features and outcomes. *Critical Care Medicine* 26:1834–1841.
- Chrusch CA, Olafson KP, McMillan PM, Roberts DE, Gray PR (2009) High occupancy increases the risk of early death or readmission after transfer from intensive care. *Critical Care Medicine* 37:2753–2758.
- de Bruin AM, van Rossum AC, Visser MC, Koole GM (2007) Modeling the emergency cardiac in-patient flow: An application of queuing theory. *Health Care Management Sci.* 10(2):125–137.
- Dobson G, Lee H-H, Pinker E (2010) A model of ICU bumping. *Oper. Res.* 58:1564–1576.
- Durbin CG, Kopel RF (1993) A case-control study of patients readmitted to the intensive care unit. *Critical Care Medicine* 21:1547–1553.
- Durrett R (1996) *Probability: Theory and Examples* (Duxbury Press, Pacific Grove, CA).
- Escobar GJ, Greene JD, Scheirer P, Gardner MN, Draper D, Kipnis P (2008) Risk-adjusting hospital inpatient mortality using automated inpatient, outpatient, and laboratory databases. *Medical Care* 46:232–239.
- Franklin C, Jackson D (1983) Discharge decision-making in a medical ICU: Characteristics of unexpected readmissions. *Critical Care Medicine* 11:61–66.
- Franklin C, Rackow EC, Mamdani B, Burke G, Weil MH (1990) Triage considerations in medical intensive care. *Arch. Internal Medicine* 150:1455–1459.
- Gajic O, Malinchoc M, Comfere TB, Harris MR, Achouiti A, Yilmaz M, Schultz MJ, Hubmayr RD, Afessa B, Farmer JC (2008) The stability and workload index for transfer score predicts unplanned intensive care unit patient readmission: Initial development and validation. *Critical Care Medicine* 36:676–682.
- Green LV (2003) How many hospital beds? *Inquiry* 39:400–412.
- Green LV (2006) Queuing analysis in healthcare. R. Hall, ed. *Patient Flow: Reducing Delay in Healthcare Delivery* (Springer, New York).
- Green LV, Savin S (2008) Reducing delays for medical appointments: A queuing approach. *Oper. Res.* 56:1526–1538.
- Green LV, Savin S, Wang B (2006a) Managing patient service in a diagnostic medical facility. *Oper. Res.* 54:11–25.
- Green LV, Soares J, Giglio JF, Green RA (2006b) Using queuing theory to increase the effectiveness of emergency department provider staffing. *Academic Emergency Medicine* 13:61–68.
- Halpern NA, Pastores SM (2010) Critical care medicine in the United States 2000–2005: An analysis of bed numbers, occupancy rates, payer mix, and costs. *Critical Care Medicine* 38:65–71.
- Huang XA (1995) A planning model for requirement of emergency beds. *J. Math. Appl. Medicine Biology* 12:345–353.
- Iwashyna TJ, Kramer AA, Kahn JM (2000) Intensive care unit occupancy and patient outcomes. *Critical Care Medicine* 37:1545–1557.
- Kc D, Terwiesch C (2012) An econometric analysis of patient flows in the cardiac intensive care unit. *Manufacturing Service Oper. Management* 14(1):50–65.
- Kwak N, Lee C (1997) A linear programming model for human resource allocation in a health-care organization. *J. Medical Systems* 21:129–140.
- Loynes RM (1963) The stability of a queue with non-independent inter-arrival and service times. *Proc. Cambridge Philos. Soc.* 58:497–530.
- McConnell KJ, Richards CF, Daya M, Weathers CC, Lowe RA (2006) Ambulance diversion and lost hospital revenues. *Ann. Emergency Medicine* 48:702–710.

- McConnell KJ, Richards CF, Daya M, Bernell SL, Weathers CC, Lowe RA (2005) Effect of increased ICU capacity on emergency department length of stay and ambulance diversion. *Ann. Emergency Medicine* 45:471–478.
- Melnick GA, Nawathe AC, Bamezai A, Green L (2004) Emergency department capacity and access in California 1990–2001: An economic analysis. *Health Affairs* 23. <http://content.healthaffairs.org/cgi/content/abstract/hlthaff.w4.136v1>.
- Merrill CT, Elixhauser A (2005) Hospitalization in the United States, 2002: *HCUP Fact Book*, No. 6 (Agency for Healthcare Research and Quality, Rockville, MD).
- Metcalf MA, Sloggett A, McPherson K (1997) Mortality among appropriately referred patients refused admission to intensive-care units. *Lancet* 350:7–11.
- Mitchell I, Grounds M, Bennett D (1995) Intensive care in the ailing UK health care system. *Lancet* 345(8950):652.
- Moreno RP, Metnitz PG, Almeida E, Jordan B, Bauer P, Campos RA, Iapichino G, Edbrooke D, Capuzzo M, Le Gall JR (2005) SAPS 3-From evaluation of the patient to evaluation of the intensive care unit. Part 2: Development of a prognostic model for hospital mortality at ICU admission. *Intensive Care Medicine* 31:1345–1355.
- Murray M, Davies M, Boushon B (2007) Panel size: How many patients can one doctor manage? *Family Practice Management* 14:44–51.
- Naylor MD, Brooten DA, Campbell RL, Maislin G, McCauley KM, Schwartz JS (2004) Transitional care of older adults hospitalized with heart failure: A randomized, controlled trial. *J. Amer. Geriatrics Soc.* 52(5):675–684.
- Norton SA, Hogan LA, Holloway RG, Temkin-Greener H, Buckley MJ, Quill TE (2007) Proactive palliative care in the medical intensive care unit: Effects on length of stay for selected high-risk patients. *Critical Care Medicine* 35:1530–1535.
- Patient Protection and Affordable Care Act (2010) Hospital readmissions reduction program. Sec. 3025. <http://www.gpo.gov/fdsys/pkg/PLAW-111pub/148/pdf/PLAW-111pub/148.pdf>.
- Rapoport J, Teres D, Lemeshow S (1996) Resource use implications of do not resuscitate orders for intensive care unit patients. *Amer. J. Respiratory Critical Care Medicine* 153:185–190.
- Renaud B, Santin A, Coma E, Camus N, Van Pelt D, Hayon J, Gurgui M, et al. (2009) Association between timing of intensive care unit admission and outcomes for emergency department patients with community-acquired pneumonia. *Critical Care Medicine* 37(11):2867–2874.
- Rincon F, Mayer SA, Rivolta J, Stillman J, Boden-Albala B, Elkind MSV, Marshall R, Chong JY (2010) Impact of delayed transfer of critically ill stroke patients from the emergency department to the neuro-ICU. *Neurocritical Care* 13:75–81.
- Rosenberg AL, Watts C (2000) Patients readmitted to ICUs: A systematic review of risk factors and outcomes. *CHEST* 118:492–502.
- Scheulen JJ, Li G, Kelen GD (2001) Impact of ambulance diversion policies in urban, suburban, and rural areas of central Maryland. *Academic Emergency Medicine* 8:1553–1557.
- Schull MJ, Vermuelen M, Slaughter G, Morrison L, Daly P (2004) Emergency department crowding and thrombolysis delays in acute myocardial infarction. *Ann. Emergency Medicine* 44:577–585.
- Shmueli A, Sprung CL, Kaplan EH (2003) Optimizing admissions to an intensive care unit. *Health Care Management Sci.* 6(3):131–136.
- Smirnov N (1939) Estimating the deviation between the empirical distribution functions of two independent samples. *Moscow Univ. Math. Bull.* 2:3–16.
- Smith GB, Taylor BL, McQuillan PJ, Nials E (1995) Rationing intensive care. Intensive care provision varies widely in Britain. *BMJ* 310:1412–1413.
- Snow N, Bergin KT, Horrigan TP (1985) Readmission of patients to the surgical intensive care unit: Patient profiles and possibilities for prevention. *Critical Care Medicine* 13:961–985.
- Swenson MD (1992) Scarcity in the intensive care unit: Principles of justice for rationing ICU beds. *Amer. J. Medicine* 92:552–555.
- Task Force of the American College of Critical Care Medicine, Society of Critical Care Medicine (1999) Guidelines for intensive care unit admission, discharge, and triage. *Critical Care Medicine* 27:633–638.
- Yankovic N, Green L (2011) Identifying good nursing levels: A queuing approach. *Oper. Res.* 59:942–955.
- Yoon KB, Koh SO, Han DW, Kang OC (2004) Discharge decision-making by intensivists on readmission to the intensive care unit. *Yonsei Med J* 45:193–198.
- Zimmerman JE, Kramer AA, McNair DS, Malila FM (2006) Acute physiology and chronic health evaluation (APACHE) IV: Hospital mortality assessment for today's critically ill patients. *Critical Care Medicine* 34:1297–1310.

Carri W. Chan is an assistant professor in the Division of Decision, Risk, and Operations at Columbia Business School. Her research interests include modeling of complex stochastic systems, efficient algorithmic design for queuing systems, and dynamic control of stochastic processing systems. Applications for this line of research include healthcare operations management, telecommunication networks, and information services.

Vivek F. Farias is the Robert N. Noyce Associate Professor of Management at the Sloan School and the Operations Research Center at the Massachusetts Institute of Technology. He works on revenue management, dynamic optimization, and the analysis of complex stochastic systems. The paper in this issue is part of the author's research in healthcare operations.

Nicholas Bambos is a professor in the Electrical Engineering Department and the Management Science and Engineering Department at Stanford University. His current research interests are in performance engineering of communication networks and computing systems, including queuing and scheduling issues in wireless and wireline networks, as well as ergodic random processes, queuing theory, and adaptive control of stochastic processing networks.

Gabriel J. Escobar is a physician and a health services researcher. He is the Regional Director for Hospital Operations Research for the Permanente Medical Group, Inc., and Director of the Systems Research Initiative at Kaiser Permanente Northern California's Division of Research. His research interests include severity of illness scoring, risk adjustment, and the use of modern electronic medical record systems for research and the improvement of patient safety.