

Methods

Allocation of Intensive Care Unit Beds in Periods of High Demand

 Huiyin Ouyang,^a Nilay Tanik Argon,^b Serhan Ziya^b
^a Faculty of Business and Economics, University of Hong Kong, Hong Kong; ^b Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina 27599

Contact: oyhy@hku.hk,  <https://orcid.org/0000-0002-0020-1860> (HO); nilay@email.unc.edu,

 <https://orcid.org/0000-0002-6814-0849> (NTA); ziya@email.unc.edu,  <https://orcid.org/0000-0003-1558-6051> (SZ)

Received: December 14, 2015

Revised: November 9, 2017; December 19, 2018

Accepted: April 3, 2019

Published Online in Articles in Advance:
 February 10, 2020

Subject Classifications: hospitals: healthcare, optimal control, dynamic programming; applications: dynamic programming

Area of Review: Stochastic Models

<https://doi.org/10.1287/opre.2019.1876>
Copyright: © 2020 INFORMS

Abstract. The objective of this paper is to use mathematical modeling and analysis to develop insights into and policies for making bed allocation decisions in an intensive care unit (ICU) of a hospital during periods when patient demand is high. We first develop a stylized mathematical model in which patients' health conditions change over time according to a Markov chain. In this model, each patient is in one of two possible health stages, one representing the critical and the other representing the highly critical health stage. The ICU has limited bed availability and therefore when a patient arrives and no beds are available, a decision needs to be made as to whether the patient should be admitted to the ICU and if so, which patient in the ICU should be transferred to the general ward. With the objective of minimizing the long-run average mortality rate, we provide analytical characterizations of the optimal policy under certain conditions. Then, based on these analytical results, we propose heuristic methods, which can be used under assumptions that are more general than what is assumed for the mathematical model. Finally, we demonstrate that the proposed heuristic methods work well by a simulation study, which relaxes some of the restrictive assumptions of the mathematical model by considering a more complex transition structure for patient health and allowing for patients to be possibly queued for admission to the ICU and readmitted from the general ward after they are discharged.

Funding: This work was supported by the U.S. National Science Foundation Division of Civil, Mechanical and Manufacturing Innovation [Grants CMMI-1234212 and CMMI-1635574].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/opre.2019.1876>.

Keywords: healthcare operations • dynamic control • Markov decision processes

1. Introduction

Efficient management of intensive care unit (ICU) beds has long been a topic of interest in practice as well as academia. Simply put, an ICU bed is a very expensive resource and the number of available ICU beds frequently falls short of the existing demand in many hospitals. Therefore, it is important to make the best use of these beds via intelligent admission and discharge decisions. There is wide agreement that during times of high demand, beds should not be given to patients who have little to benefit from intensive care treatment. However, when it comes to choosing among patients who can potentially benefit from such treatment, there do not appear to be easy answers. Even if one can quantify the ICU benefit at the individual patient level and there is agreement on some utilitarian objective such as maximizing the expected number of survivors, it is not difficult to see that allocating beds to those with the highest potential to benefit is not necessarily the right thing to do. For example, if this potential benefit can only be

realized at the expense of a long length of stay, which is likely to prevent the use of the bed for treating other patients, then it is difficult to weigh the benefits against the costs. In short, making patient admission and discharge decisions for a particular patient, especially when overall demand is high, is a complex task that requires careful consideration of not only the health condition of that particular patient in isolation but a collective assessment of the health conditions and operational requirements of all the patients in the ICU as well as the mix of patients the ICU expects to see in the near future. The objective of this paper is to use mathematical modeling and analysis to develop insights and policies that can be useful when making these complex decisions in practice, particularly under conditions where there is significantly high demand for limited ICU bed capacity.

The general framework we use to fulfill the objective we have outlined is as follows. We first develop a stylized mathematical formulation for the ICU. This relatively simple formulation (compared with the full

complexity of the actual problem) allows us to provide characterizations for the optimal policy under certain conditions. These characterizations not only provide overall insights into good ICU admit/discharge decisions, but also lead to the development of several heuristic policies that can potentially be used in practice. Finally, we test the performances of these policies with a simulation study relaxing some of the restrictive assumptions of the stylized mathematical formulation and find that the policies we propose perform quite well in comparison with some alternative benchmarks.

Our mathematical formulation assumes that each patient's health condition changes over time. Specifically, there are two discrete-time Markov chains with one representing the evolution of the patients in the ICU and one representing the evolution of the patients outside the ICU. Each Markov chain has four states corresponding to death, highly critical, critical, and survival, where death and survival states are absorbing states. Once patients enter the death state or the survival state, they leave the system, vacating the bed they had been occupying, and therefore any patient in the system can only be in one of the two health stages, critical or highly critical. In each time period, a patient arrives with some probability and a decision needs to be made as to whether to admit the patient and/or discharge any of the highly critical or critical patients to the general ward early. The objective is to minimize the long-run average number of deaths.

We start our mathematical analysis by first considering an extreme setting, where the ICU has a single bed. The main insight that comes out of this analysis is that the decision of which patient to admit to the ICU depends on how much benefit the patients are expected to get from ICU treatment and how long they are expected to stay in the ICU, and that which one of these two factors is more dominant depends on the overall level of demand for the ICU. We then consider the general setting, where the ICU has some arbitrary but finite number of beds. We formulate the decision problem as a Markov decision process (MDP) and prove that in general the optimal policy is state dependent, where the admission/discharge decisions depend on the mix of patients present in the ICU at the time the decisions are made.

Although our mathematical analysis leads to useful insights into ICU patient admit/discharge decisions, it does not directly answer the question of how one can turn these insights into practical policies and how such policies would perform under realistic conditions. To address that, we introduce a simulation model, which enriches the mathematical model in a number of directions, making it a more realistic environment for proposing and testing heuristics. Specifically, for this simulation model, we assume that patients can be in one of six health stages and they can transition from one stage to another according to a transition probability structure that is more complex than the one

assumed in the mathematical model. Unlike the case in the mathematical model, patients who have already been discharged to the general ward are also considered for readmission to the ICU, and patients who are initially admitted to the general ward can be admitted to the ICU later. Finally, in accordance with our focus on bed allocation decisions during periods of high demand, the model considers a 36-week time horizon with a 12-week period in the middle during which the ICU observes higher than usual demand levels with the arrival rate of patients first increasing and then decreasing and going back to regular levels. (The scenario is created based on the estimates of the U.S. Centers for Disease Control and Prevention (CDC) for flu seasons.) All of these additional features lead to an environment that is significantly different from the one assumed by our mathematical model. Nevertheless, the relative simplicity of our structural results make it possible for us to propose policies that can be used under more general conditions, such as those assumed by our simulation framework.

Specifically, we propose three heuristic policies and compare their performances with those of four benchmarks. The three heuristics are the ratio policy (RP), the aggregated ratio policy (ARP), and the aggregated optimal policy (AOP). RP prioritizes patients according to their expected net benefit from the ICU (increase in survival probability as a result of being treated in ICU) divided by their expected length of stay; ARP is a version of RP that assumes four patient health classes (same as assumed in the mathematical model); and AOP essentially uses the optimal policy for the mathematical model by assuming the same classifications as ARP. Note that RP and ARP are both state-independent policies whereas AOP is a state-dependent policy. Our simulation results indicate that even though all three heuristics perform well compared with the benchmarks, RP is the best policy overall. The fact that the best-performing policy is state independent suggests that the optimality of state-dependent policies established for the mathematical model may not hold in general or that it might be difficult to identify good state-dependent policies. However, it is important to note that, as we explain in detail in the paper, even though our simulation study helps us further develop our intuition into what kind of policies are likely to perform well, one should refrain from reaching definite practical conclusions mainly because research on ICU patients is not at a level where we have a clear understanding of how one should model the health condition of a patient and its evolution and as a result there is significant uncertainty as to what the right simulation model is. Therefore, one should not ignore the possibility that state-dependent policies might be superior and future studies should continue to consider them. In any case, however, the good performance of RP in our simulation study is promising for the future as it

suggests that simple policies like RP, which only requires estimates on patients' survival probability and expected length of stay, could be good enough and there may not be a need for more sophisticated decision-making tools.

2. Literature Review

In the medical literature, there has been a long line of research on quantifying the benefits of ICU care and providing empirical and mathematical support for making more sound ICU admission/discharge decisions. Most of this work has concentrated on predicting patient mortality in the ICU, estimating the benefits of ICU care, and more generally developing patient severity scores. We do not attempt to provide a thorough review of this literature here, as it is extensive and is not directly related to this paper, but only highlight a few papers as examples.

Strand and Flaatten (2008) provide a review of some of the severity scoring systems that have been proposed and used over the years. Among these scoring systems are Acute Physiology and Chronic Health Evaluation (APACHE) I, II, III, and IV (Zimmerman et al. 2006), Simplified Acute Physiology Score (SAPS) I, II, and III (Moreno et al. 2005), and Sequential Organ-Failure Assessment (SOFA) (Vincent et al. 1996). One of the objectives behind the development of these scoring systems is to obtain a tool that can reliably predict patient mortality, which has been the subject of many other articles that aimed to improve upon the predictive power of the proposed scoring systems (see, e.g., Rocker et al. 2004, Gortzis et al. 2008, and Ghassemi et al. 2014).

A number of papers study the benefits of ICU care and the effects of rationing beds in times of limited availability. Sinuff et al. (2004) review past studies on bed rationing and find that admission to the ICU is associated with lower mortality. Shmueli and Sprung (2005) study the potential survival benefit for patients of different types and severity (measured by APACHE II score) and Kim et al. (2014) quantify the cost of ICU admission denial on a number of patient outcomes including mortality, readmission rate, and hospital length of stay using a large data set. Kim et al. (2014) also carry out a simulation study to test various patient admission policies and find that a threshold-type policy that takes into account the patient severity and ICU occupancy level has the potential to significantly improve overall performance.

Studies found that delayed admission to or early discharge from ICUs, which are both common, affect patient outcomes. For example, Chalfin et al. (2007) and Cardoso et al. (2011) study patients immediately admitted to the ICU and those who had delayed admissions (i.e., waited longer than six hours for admission) and conclude that the patients in the latter

group are associated with longer length of stay and higher ICU and hospital mortality. Wagner et al. (2013) and Kc and Terwiesch (2012) find patients are discharged more quickly when ICU occupancy is high, and such patients are associated with increased mortality rate and readmission probability.

In addition to Kim et al. (2014), which we have already mentioned, a number of papers from the operations literature develop and analyze models with the goal of generating insights into capacity related questions for ICUs and step-down units (SDUs) and how patient admission and discharge decisions should be made. Modeling the ICU as an M/M/c/c queue, Shmueli et al. (2003) compare three patient admission policies and find that restricting admission to those whose expected benefit is above a certain threshold (which may or may not depend on the number of occupied beds in the ICU) brings sizeable improvements in the expected number of survivors. Dobson et al. (2010), on the other hand, develop a model in which patients are bumped out of (early discharged from) the ICU and show how this model can be used to predict performance measures like the probability of being bumped for a randomly chosen patient. The model assumes that each patient's length of stay can be observed upon arrival and when a patient needs to be bumped because of lack of beds, the patient with the shortest remaining length of stay is bumped out of the ICU. Chan et al. (2014) develop a fluid formulation in which service rate can be increased (which can be seen as patient early discharge) at the expense of increased probability of readmission. The authors identify scenarios under which taking such action is and is not helpful. Armony et al. (2018) develop a queueing model for an ICU together with an SDU and using this model provide insights into the optimal size for the SDU.

To our knowledge, within the operations literature on ICUs, the paper that is closest to our work is Chan et al. (2012). The authors consider a discrete-time MDP in which a decision needs to be made as to which patient to early discharge (with a cost) every time a new patient arrives for admission to the ICU. They show that the greedy policy, which discharges the class with the smallest discharge cost, is optimal when patient types can be ordered so that the types with smaller discharge costs have shorter expected length of stay and provide bounds on the performance of this policy for cases when such ordering is not possible. Despite some similarities, our formulation and analysis have some important differences. We assume that patients can be in one of two health stages, can transition from one stage to the other during their stay, and eventually either die or survive. On the other hand, Chan et al. (2012) allow for multiple types of patients whose health status can also

change over time, but their model does not permit a patient to return to a previous state. The main reason why these differences are important is that the analysis of the two models leads to two different sets of results which complement each other. In particular, our formulation allows us to push the analytical results and optimal policy characterizations further and thereby provide deeper insights into optimal ICU admission and discharge decisions. For example, we provide a characterization of the optimal policy not only when patients with higher benefits from ICU have shorter length of stay but also when higher benefits can only come at the expense of longer length of stay in the ICU.

Our analysis in this paper can also be seen as a contribution to the classical queueing control literature where arriving jobs are admitted or rejected according to some reward or cost criteria. More specifically, because jobs in our model do not queue, it can be seen as a loss system (see, e.g., Örmeci et al. 2001, Örmeci and Burnetas 2005, Ulukus et al. 2011 and references therein). Within this literature, Ulukus et al. (2011) appears to be the closest to our work. This paper considers a model in which the decision is not only whether an arriving job should be admitted but also whether any of the jobs in service should be terminated. This termination action can be seen as the early discharge action in our model. However, despite this similarity, there are some important differences in the formulation. Whereas Ulukus et al. (2011) consider a more general form for the termination cost and multiple job classes, they do not allow the possibility of jobs changing types during service. There are also important differences in the results. Just as we do in this paper, Ulukus et al. (2011) also provide conditions under which one of the two types should be preferred over the other at all times. However, our formulation makes it possible for us to provide optimal policy characterizations at a more detailed level and mathematically establish some of the numerical observations made by Ulukus et al. (2011) regarding the threshold structure of the optimal policy.

3. Model Description

In this section, we describe the stylized mathematical formulation we use to generate insights into good bed allocation decisions and develop practical heuristic methods. Specifically, in this model, we consider an ICU with a capacity of b beds, where b is a finite positive integer. Patients arriving to this system are assumed to have health conditions that require treatment in an ICU. However, there is also the option of admitting these patients to what we refer to as the general ward, where the patient may be provided a different level of service. It is also possible that a

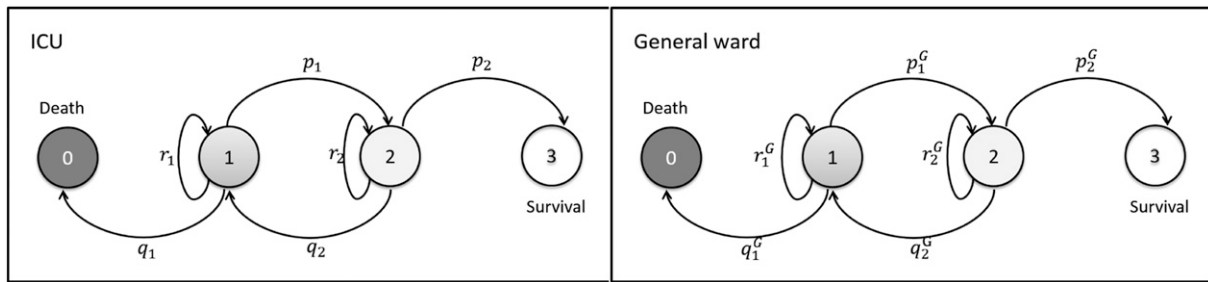
patient who was previously admitted to the ICU can be early discharged to the general ward to accommodate another patient. Note that we use the general ward to represent any non-ICU care unit, which includes actual hospital wards, step-down or transitional care units, nursing homes, and any other facility that can accommodate the patients but cannot provide an ICU-level service to the patients. In our model, we assume that all these non-ICU beds are identical and the capacity of the general ward is infinite.

Arriving patients are assumed to be in one of two health stages, with stage 1 representing a highly critical condition and stage 2 representing a critical condition. We consider discrete time periods during which at most one patient arrives. Let $\lambda_i > 0$ denote the probability that a stage i patient will arrive in each period for $i = 1, 2$; let $\lambda \equiv \lambda_1 + \lambda_2$ denote the probability that there will be a patient arrival; let $\bar{\lambda} \equiv 1 - \lambda$ denote the probability of no arrival, where we assume $\lambda < 1$. During their stay, in the ICU or in the general ward, patients' health conditions change according to a Markov chain and they eventually either enter stage 0 or stage 3. Stage 0 corresponds to the death of the patient whereas stage 3 represents the patient's survival. As soon as patients hit either stage 0 or 3, they leave the system, vacating the bed they had been occupying. We assume that the system incurs a unit cost every time a patient leaves in stage 0 whereas there is no cost or reward associated with other stages.

Patients currently in stage $i \in \{1, 2\}$ can enter stage $i + 1$ or $i - 1$ in the next time period with probabilities that depend on where they are being treated: ICU or general ward. A stage i patient in the ICU either jumps to stage $i + 1$ with probability p_i , jumps to stage $i - 1$ with probability q_i , or stays in stage i with probability $r_i = 1 - p_i - q_i$. The respective probabilities for the general ward are p_i^G, q_i^G and r_i^G . We assume that p_i, q_i, p_i^G, q_i^G are all strictly positive whereas r_i and r_i^G are nonnegative. The transition diagram of patient evolution is shown in Figure 1.

In some respects, assuming that sick patients can only be in one of two health stages can be seen as a significant simplification of reality. Although it is true that it is difficult to capture the full spectrum of patient diversity with a two-stage model, the assumption helps us capture the reality that patients' health conditions change over time at least in some stylized way without rendering the analysis impossibly difficult. More importantly, the assumption can in fact be justified in some contexts because even in practice such simplifications are made to bring highly complex decision problems to manageable levels. When managing patient demand under highly resource-restrictive environments, particularly in case of epidemics and mass-casualty events, practitioners typically choose to employ prioritization policies that

Figure 1. Transition Diagram of Patient Evolution in the ICU and General Ward



keep the number of triage classes at minimum in an effort to make the policies simpler and easier to implement. For example, the ICU triage protocol developed by Christian et al. (2006) places patients in need of ICU treatment into one of two priority classes based on the patients’ SOFA scores. The proposed protocol also calls for patient reassessments recognizing the possibility that there could be changes in the patients’ health conditions. Nevertheless, in Section 6, we consider a more detailed and arguably more realistic evolution model for patients’ health condition and demonstrate how our analysis based on this rather simplified structure would be useful.

At each time period, the decision maker needs to make the following decisions: (i) if there is an arrival, whether the patient should be admitted to the ICU or the general ward, and (ii) which patients in the ICU (if any) should be early discharged to the general ward regardless of whether there is a new arrival or not. Note that if all b beds are occupied at the time a stage i patient arrives, admitting the patient will mean early discharging at least one stage $3 - i$ patient to the general ward. To keep the presentation simple, we will call both the decision of discharging an existing patient from the ICU to the general ward and admitting a new arrival to the general ward “discharge” even though the latter action does not in fact correspond to a discharge but direct admission to the general ward.

We formulate this problem as an MDP. We denote the system state by $\mathbf{x} = (x_1, x_2)$, where x_i represents the number of stage i patients. Note that any new arrival is included either in x_1 or x_2 since there is no need to distinguish between new and existing patients. Since the ICU has a capacity of b and at most one patient arrives in each time period, the state space is

$$\mathcal{S} = \{(x_1, x_2) : x_1, x_2 \geq 0 \text{ and } x_1 + x_2 \leq b + 1\}.$$

The decision at each epoch can be described by action $\mathbf{a} = (a_1, a_2)$, where a_i is the number of stage i patients to be discharged. The action space is defined as

$\mathcal{A} = \{(a_1, a_2) : a_1, a_2 \geq 0, \text{ and } a_1 + a_2 \leq b + 1\}$. Then in any state $(x_1, x_2) \in \mathcal{S}$, the feasible action set is

$$\mathcal{A}(x_1, x_2) = \{(a_1, a_2) : 0 \leq a_i \leq x_i, \text{ for } i = 1, 2, \text{ and } x_1 + x_2 - a_1 - a_2 \leq b\}.$$

Let ϕ_i^G denote the probability that a patient who is discharged to the general ward in stage i will end up in stage 0 for $i = 1, 2$. Then, ϕ_i^G can be computed by solving the following equations:

$$\phi_1^G = q_1^G + r_1^G \phi_1^G + p_1^G \phi_2^G, \quad \phi_2^G = q_2^G \phi_1^G + r_2^G \phi_2^G.$$

Letting $\beta_i^G = q_i^G / p_i^G$ for $i = 1, 2$, we can show that

$$\phi_1^G = \frac{\beta_1^G + \beta_1^G \beta_2^G}{1 + \beta_1^G + \beta_1^G \beta_2^G}, \quad \phi_2^G = \frac{\beta_1^G \beta_2^G}{1 + \beta_1^G + \beta_1^G \beta_2^G}. \quad (1)$$

Similarly, for $i = 1, 2$, let ϕ_i denote the probability that a patient who is admitted to the ICU in stage i will end up in stage 0 under the condition that the patient will never be early discharged to the general ward. Then, ϕ_i can similarly be computed as

$$\phi_1 = \frac{\beta_1 + \beta_1 \beta_2}{1 + \beta_1 + \beta_1 \beta_2}, \quad \phi_2 = \frac{\beta_1 \beta_2}{1 + \beta_1 + \beta_1 \beta_2}, \quad (2)$$

where $\beta_i = q_i / p_i$ for $i = 1, 2$.

Let $c(x_1, x_2, a_1, a_2)$ denote the immediate expected cost of taking action (a_1, a_2) in state (x_1, x_2) . The expected cost for the patients who will occupy the ICU during the next period is equal to the expected number of ICU patients who will transition to state 0 in the next time period; that is, $(x_1 - a_1)q_1$. The expected cost for the discharged stage i patients is $a_i \phi_i^G$ since each discharged patient will end up in stage 0 with probability ϕ_i^G . Note that this second portion of the cost is the expected lump-sum cost of discharging stage i patients, the expected cost that will eventually incur, not the immediate cost. However, for our analysis, we can equivalently assume that this cost will incur immediately since we know that if the patient enters state 0 eventually, this will happen

within some finite time period with probability 1. The total immediate expected cost then can be written as

$$c(x_1, x_2, a_1, a_2) = a_1\phi_1^G + a_2\phi_2^G + q_1(x_1 - a_1).$$

Note that although a hospital could possibly also have financial considerations when making patient admit/discharge decisions, particularly under non-emergency conditions, in this paper, in parallel with our focus on periods during which there is excessively high demand, we restrict our focus to policies that aim to minimize the number of deaths.

Let $P_{(a_1, a_2)}(x_1, x_2, y_1, y_2)$ denote the probability that the system will transition to state (y_1, y_2) from state (x_1, x_2) when action (a_1, a_2) is chosen. Then, we have $P_{(a_1, a_2)}(x_1, x_2, y_1, y_2) = P(y_1, y_2 | x_1 - a_1, x_2 - a_2)$, where $P(y_1, y_2 | x_1, x_2)$ denotes the probability that given that there are x_1 stage 1 patients and x_2 stage 2 patients at a decision epoch after that epoch's action is taken, there will be y_1 stage 1 patients and y_2 stage 2 patients at the beginning of the next decision epoch. Specifically,

$$\begin{aligned} P(y_1, y_2 | x_1, x_2) &= \bar{\lambda} \sum_{u=0}^{x_1} \sum_{d=0}^{x_1-u} \bar{P}_1\{x_1, u, d\} \bar{P}_2\{x_2, x_1 + x_2 \\ &\quad - d - y_1 - y_2, y_1 - (x_1 - u - d)\} \\ &\quad + \lambda_1 \sum_{u=0}^{x_1} \sum_{d=0}^{x_1-u} \bar{P}_1\{x_1, u, d\} \bar{P}_2\{x_2, x_1 + x_2 \\ &\quad - d - (y_1 - 1) - y_2, (y_1 - 1) - (x_1 - u - d)\} \\ &\quad + \lambda_2 \sum_{u=0}^{x_1} \sum_{d=0}^{x_1-u} \bar{P}_1\{x_1, u, d\} \bar{P}_2\{x_2, x_1 + x_2 \\ &\quad - d - y_1 - (y_2 - 1), y_1 - (x_1 - u - d)\}, \end{aligned}$$

where $\bar{P}_i\{x_i, u, d\}$ is the probability that of the x_i stage i patients, u of them will transition to stage $i + 1$ and d of them will transition to stage $i - 1$; that is,

$$\bar{P}_i\{x_i, u, d\} = \begin{cases} \binom{x_i}{u} \binom{x_i-u}{d} p_i^u q_i^d r_i^{x_i-u-d}, & \text{for } u, d \geq 0 \\ & \text{and } u + d \leq x_i \\ 0, & \text{otherwise.} \end{cases}$$

A policy π maps the state space \mathcal{S} to the action space \mathcal{A} . We use Π to denote the set of feasible stationary discharge policies. Let $N_\pi(t)$ and $N_\pi^G(t)$, respectively, denote the number of patients who enter stage 0 by time t in the ICU and in the general ward. Then $J^\pi(\mathbf{x})$, the expected long-run average cost under policy π given the initial state \mathbf{x} , can be expressed as

$$J^\pi(\mathbf{x}) = \lim_{t \rightarrow \infty} \frac{1}{t} E[N_\pi(t) + N_\pi^G(t) | \mathbf{x}].$$

Our objective is to obtain an optimal policy π^* such that $J^{\pi^*}(\mathbf{x}) \leq J^\pi(\mathbf{x})$ for any $\pi \in \Pi$ and $\mathbf{x} \in \mathcal{S}$. Note that this MDP is a unichain with finite state and action

spaces, hence the above limit exists and is independent of the initial state \mathbf{x} (see, e.g., theorem 8.4.5 of Puterman 2005). We also know that there exists a bounded function $h(x_1, x_2)$ for $(x_1, x_2) \in \mathcal{S}$ and a constant g satisfying the optimality equation

$$\begin{aligned} h(x_1, x_2) + g = \min_{(a_1, a_2) \in \mathcal{A}(x_1, x_2)} &\left\{ c(x_1, x_2, a_1, a_2) \right. \\ &\left. + \sum_{(y_1, y_2) \in \mathcal{S}} P_{(a_1, a_2)}(x_1, x_2, y_1, y_2) h(y_1, y_2) \right\}, \end{aligned} \tag{3}$$

and there exists an optimal stationary policy π^* such that $g = J^{\pi^*}(\mathbf{x})$ and π^* chooses an action that maximizes the right-hand side of (3) for each $(x_1, x_2) \in \mathcal{S}$.

4. Single-Bed ICU

In this section, we consider the case where $b = 1$, that is, there is a single ICU bed. The objective of this analysis is to generate insights into situations where ICU capacity is severely limited. It will also provide support for one of the heuristic policies we propose in Section 6.

When $b = 1$, at any decision epoch there are at most two patients under consideration, the patient who is currently occupying the bed (if there is one) and the patient who has just arrived for possible admission (if there is an arrival). Restricting ourselves to nonidling policies, (i.e., the bed is never left empty when there is demand), we investigate the question of which of the two patients to admit to the ICU. (An implicit assumption here is that the ICU is the preferred environment for the patients. This is a reasonable assumption to make, but nevertheless in the next section, we identify conditions under which this is true in our mathematical formulation.) Specifically, there are two stationary policies to compare, $\bar{\pi}_1$, the policy that discharges the stage 1 patient, and $\bar{\pi}_2$, the policy that discharges the stage 2 patient when the choice is between a stage 1 and a stage 2 patient. Under any of the two policies, when there are two patients in the same stage, the choice between the two is arbitrary. Let $J^{\bar{\pi}_k}$ for $k \in \{1, 2\}$ denote the long-run average cost under policy $\bar{\pi}_k$.

The following proposition provides a comparison of the performances of the two policies, which accounts for both the incremental survival benefit and the required ICU length of stay when making prioritization decisions. (The proof for the proposition as well as the proofs of all the other analytical results in the paper are provided in the online appendix.) We first let L_i denote the expected ICU length of stay for a patient admitted to the ICU in stage i and is never early discharged in either

stage 1 or 2. Then, L_i can be obtained by solving the equations $L_1 = 1 + r_1L_1 + p_1L_2$ and $L_2 = 1 + q_2L_1 + r_2L_2$, which gives us

$$L_1 = \frac{p_1 + p_2 + q_2}{p_1p_2 + q_1p_2 + q_1q_2}, L_2 = \frac{p_1 + q_1 + q_2}{p_1p_2 + q_1p_2 + q_1q_2}. \quad (4)$$

Proposition 1. Suppose that $b = 1$, that is, there is a single ICU bed, and the ICU admission decision is between a stage 1 and stage 2 patient. Also assume without loss of generality that $\phi_i^G - \phi_i \geq \phi_{3-i}^G - \phi_{3-i}$ for some fixed $i \in \{1, 2\}$. Then, we have

(a) if $\frac{\phi_i^G - \phi_i}{L_i} \geq \frac{\phi_{3-i}^G - \phi_{3-i}}{L_{3-i}}$, then it is optimal to admit the patient in stage i , that is, $J^{\bar{\pi}_i} \geq J^{\bar{\pi}_{3-i}}$;

(b) if $\frac{\phi_i^G - \phi_i}{L_i} < \frac{\phi_{3-i}^G - \phi_{3-i}}{L_{3-i}}$, then it is optimal to admit the patient in stage i , that is, $J^{\bar{\pi}_i} \geq J^{\bar{\pi}_{3-i}}$, if and only if

$$\lambda \leq \frac{(\phi_i^G - \phi_i) - (\phi_{3-i}^G - \phi_{3-i})}{(\phi_i^G - \phi_i) - (\phi_{3-i}^G - \phi_{3-i}) + [L_i(\phi_{3-i}^G - \phi_{3-i}) - L_{3-i}(\phi_i^G - \phi_i)]}. \quad (5)$$

The difference $\phi_i^G - \phi_i$ can be seen as the benefit of staying in the ICU instead of the general ward for a stage i patient. From a system optimization point of view, we can call the patients with larger $\phi_i^G - \phi_i$ high-value patients. On the other hand, the ratio $(\phi_i^G - \phi_i)/L_i$ can roughly be seen as the per unit time benefit of keeping a patient who arrives in stage i in the ICU at all times and thus we can call the patients with larger $(\phi_i^G - \phi_i)/L_i$ high-value-rate patients. Then, according to Proposition 1(a), if stage i patients are both high-value and high-value-rate patients, they should be preferred over stage $3 - i$ patients. As Proposition 1(b) implies, for stage i patients to be preferable, it is not sufficient for them to be high value. If they are high value but not high value rate, then they are preferable only if the arrival rate is sufficiently small. This is because when the arrival rate is small, having a limited bed capacity is less of a concern and thus in that case the value is the dominating factor. However, when the arrival rate is large, the lengths of stay are important as they would be a key factor in the availability of the ICU beds for new patients. As a result, the rate with which the value is generated becomes the dominant factor.

These results point to the importance of taking into account the ICU load when making patient admission/early discharge decisions and prioritizing one patient over the other. In short, what may be the right thing to do for one ICU may not be right for another. For ICUs with relatively ample capacity, it might be best to focus on identifying patients who will benefit most from ICU care and admit them without being overly concerned about how long they will stay. However, for highly loaded ICUs, the decision is more complicated and

the anticipated length of stay should be part of the decision. In the following section, we investigate this question further by analyzing dynamic decisions in a model where the number of beds in the ICU can take any finite value.

5. Analysis of the Multibed ICU Model

In this section, we consider the long-run average-cost optimization problem with optimality equations given in Equation (3). An optimal action in any particular state is the one that achieves the minimum in the optimality equation. We denote the set of optimal actions in state (x_1, x_2) by $\mathcal{A}^*(x_1, x_2)$:

$$\begin{aligned} \mathcal{A}^*(x_1, x_2) &= \left\{ (\bar{a}_1, \bar{a}_2) \in \mathcal{A}(x_1, x_2) : c(x_1, x_2, \bar{a}_1, \bar{a}_2) \right. \\ &\quad \left. + \sum_{(y_1, y_2) \in \mathcal{S}} P_{(\bar{a}_1, \bar{a}_2)}(x_1, x_2, y_1, y_2) h(y_1, y_2) \right. \\ &= \min_{(a_1, a_2) \in \mathcal{A}(x_1, x_2)} \left\{ c(x_1, x_2, a_1, a_2) \right. \\ &\quad \left. + \sum_{(y_1, y_2) \in \mathcal{S}} P_{(a_1, a_2)}(x_1, x_2, y_1, y_2) h(y_1, y_2) \right\} \end{aligned}$$

Since the state space and action space are finite and costs are bounded, \mathcal{A}^* is nonempty. In general, the set $\mathcal{A}^*(x_1, x_2)$ can have more than one element. However, for convenience, we adopt the following convention for picking one action from the set and refer to it as the optimal action for state (x_1, x_2) . Specifically, we define the optimal action $a^*(x_1, x_2) = (a_1^*(x_1, x_2), a_2^*(x_1, x_2))$, where

$$a_1^*(x_1, x_2) = \min\{\bar{a}_1 : (\bar{a}_1, \bar{a}_2) \in \mathcal{A}^*(x_1, x_2)\}, \text{ and} \\ a_2^*(x_1, x_2) = \min\{\bar{a}_2 : (a_1^*(x_1, x_2), \bar{a}_2) \in \mathcal{A}^*(x_1, x_2)\}.$$

Thus, if there are multiple actions for any given state, we choose the one that discharges as few stage 1 patients as possible; if there are multiple such actions, then among those we choose the one that discharges as few stage 2 patients as possible.

Theorems 1, 2, and 3 presented in this section characterize the structure of the optimal policy. The proofs of these theorems are provided in the online appendix, where we first analyze the system with the objective of minimizing expected total discounted cost and establish some analytical properties, which serve as a stepping stone to our main results for the long-run average case.

5.1. Optimality of Nonidling ICU Beds

The nonidling policies are defined as the policies that will always allocate an ICU bed to a new arriving patient and never discharge an ICU patient to the general ward when there are ICU beds available.

We first identify conditions under which there exists an optimal policy, which is nonidling.

Theorem 1. *Suppose that $\beta_i < \beta_i^G$ for $i = 1, 2$. Then, there exists a stationary average-cost optimal policy, which is nonidling, that is, a policy under which it is never optimal to leave an ICU bed empty whenever there is a patient in need of treatment.*

Comparing β_i with β_i^G can be seen as one way of assessing the potential benefit of ICU over the general ward for stage i patients. The condition $\beta_i < \beta_i^G$ for $i = 1, 2$ essentially means that the ratio of the probability of a patient getting worse to the probability of a patient getting better over the next time step is smaller in the ICU for all the patients. Theorem 1 states that this condition is sufficient to ensure the existence of an optimal policy that admits patients of either stage to the ICU as long as there is an available bed. We found numerical examples that show that when this condition does not hold, the optimal policy is not necessarily nonidling, meaning that the ICU would only accept patients from a particular stage and keep some of the ICU beds empty even when there is demand from patients of the other stage.

5.2. General Structure of the Optimal Policy

Since we restrict ourselves to the set of nonidling policies, which we know contains an optimal policy under the assumption that $\beta_i < \beta_i^G$ for $i = 1, 2$, we only need to investigate the optimal actions for states (x_1, x_2) such that $x_1 + x_2 = b + 1$ and $x_1, x_2 > 0$, that is, when all ICU beds are currently occupied, a patient has just arrived, and there are patients from both stages (including the patient who has just arrived). As we describe in the following theorem, it turns out that the optimal decision has a threshold structure.

Theorem 2. *Suppose that $\beta_i < \beta_i^G$ for $i = 1, 2$. Then, there exists a threshold $x^* \in [1, b + 1]$ such that for any state (x_1, x_2) with $x_1, x_2 > 0$ and $x_1 + x_2 = b + 1$, we have*

$$a^*(x_1, x_2) = \begin{cases} (1, 0) & \text{if } x_1 \geq x^* \\ (0, 1) & \text{if } x_1 < x^*. \end{cases}$$

According to Theorem 2, when the nonidling condition holds and when the system conditions are so that one of the patients has to be admitted to the general ward because of a fully occupied ICU, whether that patient should be a stage 1 or stage 2 patient depends on the health conditions of all the patients in the ICU. Specifically, if the number of stage 1 (stage 2) patients in the ICU is above a particular threshold value, which depends on all the model parameter values and thus survival probabilities as well as lengths of stay, then one of the stage 1 (stage 2) patients should be admitted to the general ward.

In other words, if there are sufficiently many stage 1 patients, the preference should be for a stage 2 patient; otherwise the preference should be for a stage 1 patient.

It is important to note that although x^* can take one of the boundary values of 1 or $b + 1$ (both of which would imply that the policy is in fact not dependent on the composition of the patients), there are examples that show that it can also take values in between. This means that there are indeed certain settings in which the optimal policy is state dependent. (We should note, however, that it is not clear whether the potential benefits of using such a state-dependent policy can be realized in practice. We investigate and discuss this issue in detail in Section 6.1.)

The fact that in general the optimal policy can be state dependent might seem somewhat surprising at first because the implication is that if there are two specific patients, A and B, one of them being in stage 1 the other in stage 2, and only one of them can be admitted to the ICU, then whether we choose A or B depends on the health stages of all the patients in the ICU, not just A and B. Given that this decision will not impact other patients' survival chances and the survival chances of patients A and B do not depend on the other patients in the ICU, it is unclear why the choice between A and B depends on the other patients. To clarify this, in light of our analysis of the single-bed case, consider the two important factors that go into the decision of which patient to admit: expected net ICU benefit, which we would like to be as high as possible, and expected length of stay, which we would like to be as small as possible. The expected length of stay is important because it directly affects the bed availability for future patients. In particular, it affects the probability that a bed will be available the next time there is a patient seeking admission to the ICU. However, whether a bed will be available for the next patient (and patients thereafter) depends on the length of stay for not just patient A and patient B, but for all the patients in the ICU.

Now, consider two extreme cases, one in which patients other than A and B all have very short expected lengths of stay and one in which they all have long expected lengths of stay. In the former case, there is a good chance for a bed to be available soon even if we ignore A and B, and this, when choosing between A and B, will make the expected lengths of stay for A and B far less important compared with the latter case. Thus, in the former case, whoever has the larger expected benefit, will be (most likely) admitted to the ICU. In the latter case, however, the decision is more complicated and to make a bed available for the next patient with a higher probability, it might actually be preferable to admit the patient with the smaller expected net benefit if that patient's expected length of stay is shorter. In general, one can then see that, as the

composition of the patients in the ICU changes, future bed availability probability changes and this in turn results in shifting preferences for the patient to be admitted. More specifically, as Theorem 2 implies, there is an ideal mix of patients (a certain number of stage 1 patients and a certain number of stage 2 patients), which hits the right balance between the expected benefit and the future bed availability, and the optimal policy continuously strives to push the system to that level by employing a threshold-type policy.

Given this explanation, it would be reasonable to expect that patient A should always be preferred over patient B regardless of the patient composition in the ICU if the expected benefit for patient A is larger than that of patient B and the expected length of stay for patient A is smaller than that of patient B. We can indeed prove that is the case as we formally state in the following theorem.

Theorem 3. *Suppose that $\beta_i < \beta_i^G$ for $i = 1, 2$, and for some fixed $k \in \{1, 2\}$,*

$$\phi_k^G - \phi_k < \phi_{3-k}^G - \phi_{3-k} \text{ and } L_k \geq L_{3-k}. \quad (6)$$

Then, for any state (x_1, x_2) such that $x_1 + x_2 = b + 1$, we have $a^(x_1, x_2) = (a_1^*, a_2^*)$ with $a_k^* = 1$ and $a_{3-k}^* = 0$.*

Theorem 3 states that if a particular health stage is associated with a lower expected ICU benefit and longer expected length of ICU stay, then a patient from that health stage should be admitted to the general ward when the demand for the ICU exceeds the ICU bed capacity. In this case, the optimal policy is simple since one of the two stages can be designated as the higher priority stage regardless of the system state. The result makes sense intuitively. If patient A will benefit more from the ICU bed than patient B and patient A will also vacate the bed more quickly for the use of the future patients, there is no reason why the bed should be given to patient B.

6. Simulation Study

In Sections 4 and 5, we analyzed relatively simple formulations with the objective of generating insights and coming up with heuristic methods, which are flexible enough to be used under more general and realistic conditions. In this section, we have two main goals. First, to demonstrate how one can construct heuristic policies based on our analysis assuming that we know how the health status of patients evolve in the ICU and in the general ward, and propose specific policies for the assumed evolution model. Second, to report the findings of our simulation study where we investigated how the policies we generated perform. Our simulation model relaxes some of the restrictive assumptions of the mathematical model of Section 3. In particular, we consider a more detailed and

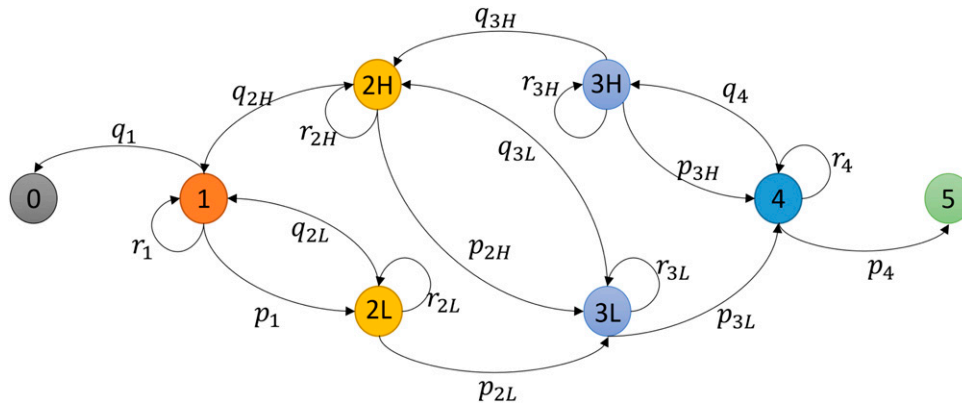
realistic health evolution model, a nonstationary patient arrival process, the possibility of patients to wait for admission to the ICU, and possible readmission of patients who have already been discharged from the ICU. We start with describing the health evolution model used in our simulation model.

6.1. Simulation Model

Given what is known in the medical literature, it is not possible to construct a detailed, realistic model for describing how each patient's health status evolves in and outside the ICU. This obviously poses a significant challenge in reaching the two goals outlined previously. Although our mathematical model, which assumes two health stages and possible transitions between the two, broadly captures what happens in practice and is in fact in line with the only proposed classification protocol developed (see Christian et al. 2006), it is also very likely that the model, with its mathematically convenient construction like having Markovian transition probabilities, fails to capture some of the features that one might see in reality. For example, two patients might be in the same health stage with respect to some objective criterion (one can think of a classification based on the SOFA score as used by Christian et al. 2006), but assuming they would have the same stochastic evolution in the future could be an oversimplification if, for instance, one of the patients has just arrived and the other patient has been in the ICU for hours, or one of the patients' health status has been gradually improving suggesting a positive trend whereas the other's health has been declining. Thus, there are a number of ways our mathematical model can be generalized to make it more realistic.

The health evolution model we used in our simulation study, which is depicted in Figure 2, helped us capture some of the features we described. The model assumes four levels of criticality, but also takes into account the direction of the last transition for the intermediate two criticality levels. Explicit consideration of the last transition makes it possible to at least partially capture the effect of trend in the evolution of patients' health status. More specifically, we assume that at any point in time patients in the ICU or the general ward belong to one of the six stages $\{1, 2H, 2L, 3H, 3L, 4\}$, where stages 1 and 4 denote the most critical and the least critical levels, stages 2H and 3H denote the intermediate criticality levels for patients whose health condition has been declining (i.e., the last transition was from a healthier stage) and stages 2L and 3L denote the intermediate criticality levels for patients whose health condition has been improving (i.e., the last transition was from a less healthy stage). Within one time period, which is assumed to be

Figure 2. (Color online) Transition Diagram for Patient Evolution in the ICU



Note. Patients in the general ward follow the same transition model with corresponding transition probabilities indicated by the superscript G.

one hour (and thus one day consists of 24 time periods), the health condition of patients in stage $i \in \{1, 2H, 2L, 3H, 3L, 4\}$ can improve with probability p_i , decline with probability q_i , or stay the same with probabilities $r_i = 1 - p_i - q_i$. Stages 0 and 5 are two absorbing stages where 0 corresponds to death and 5 corresponds to survival. (See Figure 2 to see the probabilities corresponding to each transition.)

When choosing the values for transition probabilities, rather than setting them completely randomly, we set them in a way that the system at least conforms to what we know from the medical literature. Several articles in the literature provide estimates on ICU length of stay and survival probabilities. However, in line with our focus on situations where the ICU experiences an extremely high demand over a long period of time, we chose to use the estimates that are provided by Kumar et al. (2009), which are based on data obtained in Canada during the 2009 H1N1 influenza outbreak. Kumar et al. (2009) found that the average mortality rate in the ICU was approximately 17% and the average length of stay in the ICU was 12 days. Therefore, we randomly generated scenarios so that the expected ICU death probability over all the scenarios is approximately 0.17 and the expected length of stay (with no early discharge) for the same is approximately $24 \times 12 = 288$ hours. In addition, we ensured that the generated scenarios satisfied the condition that patients who were previously in healthier stages are more likely to get better.

When generating the random scenarios, we first identified a baseline setting that conforms to the description given and then made random choices around this baseline. More specifically, we set $p_1 = 0.016$, $p_{2L} = 0.032U_{2L}$, $p_{2H} = 0.032U_{2H}$, $p_{3L} = 0.016U_{3L}$, $p_{3H} = 0.016U_{3H}$, $p_4 = 0.012$, and $q_1 = 0.0072$, $q_{2L} = 0.01V_{2L}$, $q_{2H} = 0.01V_{2H}$, $q_{3L} = 0.012V_{3L}$, $q_{3H} = 0.012V_{3H}$, $q_4 = 0.016$, where U_{2L} , U_{3L} , V_{2H} , V_{3H} are independent random variables each uniformly distributed over $(0.5, 1)$, and

V_{2L} , V_{3L} , U_{2H} , U_{3H} are independent random variables each uniformly distributed over $(1, 1.5)$. (Note that the baseline level corresponds to the case where each random variable is set to 1.)

Kumar et al. (2009) do not provide any estimates on what the survival probabilities for the ICU patients would be if they were treated outside the ICU. In the absence of such estimates, recognizing that the condition of patients treated in non-ICU wards would be more likely to become worse and less likely to become better, for each $i \in \{1, 2H, 2L, 3H, 3L, 4\}$, we obtained q_i^G by multiplying q_i by a random coefficient uniformly distributed over $(1, 2)$, and p_i^G by multiplying p_i by a random coefficient uniformly distributed over $(0.5, 1)$.

In the simulation study, we focused on a time period during which the hospital experiences the flu season. To model patient arrivals realistically, we used CDC flu season reports as well as FluSurge 2.0, the influenza patient demand prediction tool developed by CDC. As one can observe from Figure A2 in the online appendix, the flu season typically starts with a period where the arrival rate is mostly stationary, which is followed by an outbreak period, and ends with another stationary period. We considered a 36-week time period where during the first 12 weeks and the last 12 weeks patient demand is stationary (with an arrival probability of λ_{st} in each time period), and the outbreak and the nonstationary demand period is observed during the middle 12 weeks. According to the default scenario assumed by FluSurge 2.0, in this middle 12-week period, the daily percentage change in demand (DPCD) (i.e., percentage change in the expected number of new patient arrivals) is 3% during the first six weeks and -3% during the next six weeks. In our study, we considered two settings, one with DPCD value of 3%, and the other with 5% (six weeks of increase followed by six weeks of decrease with the same absolute value for the rate). For the baseline stationary arrival rate,

which the ICU observes during the first 12 weeks and the last 12 weeks, we considered three different levels. Specifically, we let the ICU load $\rho_{st}^A = \lambda_{st}E[L]/b$ (where b is the number of ICU beds) to be either 0.5, 0.8, or 1. The choice of baseline load on the ICU also determines the overall demand level during the outbreak period since the arrival rates of patients will increase starting from these baseline levels. As for the health stages for the new patients, rather than assuming that they all come in a given state, we assumed that there is patient heterogeneity. Specifically, letting θ_i denote the probability that the initial health stage for an incoming random patient is i , when generating scenarios, we let $\theta_i = (U_i^A + 1) / \sum_{j \in \{1,2L,2H,3L,3H,4\}} (U_j^A + 1)$, where $U_i^A \sim U(0,1)$ for each stage i .

We assumed that the ICU has 20 beds for our simulation study. (Note that the choice of a 20-bed ICU together with the three load levels we consider in our study are consistent with the range of possible demand predictions of FluSurge 2.0 and a typical population/ICU bed ratio in the United States.) We also assumed that, as in the mathematical model, there is no limit on the number of patients who can be accommodated in the general ward. (Note that whereas general ward beds are also limited in numbers in reality, they are more widely available than ICU beds and the key issue typically is the effective management of ICU beds.) However, in the simulation model, in accordance with commonly observed practice, we assumed that when a bed becomes available in the ICU and there are patients in the general ward, one of those patients is admitted to the newly vacated bed. With this feature, the simulation model allows the possibility of readmitting patients who were previously discharged from the ICU to the general ward back to the ICU and having patients who find the ICU full to queue up in the general ward for possible admission later on.

6.2. Proposed Policies and Benchmarks

In this section, we propose policies that are based on our mathematical analysis but are meant to be used in the more general construction assumed in the simulation model. By doing that, we will also be illustrating more generally in what way our mathematical results and insights can be used to develop heuristics that can be used under any future patient health evolution model that is supported by medical research and data. It is important to note that the policies we propose assume that patient demand is so high that ICU admits and discharges are done throughout the day as needed unlike some of the common practices in place under regular operating conditions, which restrict such decisions to be made and actions to be taken only during certain times of the day.

The first two policies described next are included mainly because they can serve as benchmark policies and do not necessarily represent policies that are used in practice.

First-Come-First-Served (FCFS). Patients are admitted to the ICU beds in the order they arrive. None of the patients are discharged early to the general ward when a new patient finds the ICU full. In such a case, the patient is admitted to the general ward and waits for an opening in the ICU. When a patient in the ICU leaves (as a result of death or survival), among the patients who are still in the general ward, the one who was first admitted to the general ward is admitted to the newly vacated bed in the ICU. This policy would clearly capture the policy of not being proactive about making the best possible use of the ICU and opting for a policy, which could largely be considered as fair rather than aiming to maximize the greatest good for the greatest number.

Random Discharge Policy (RDP). Under this policy, if the ICU is fully occupied when a patient arrives, one of the patients among the patients already in the ICU and the patient who has just arrived is randomly chosen and transferred to the general ward. When a patient in the ICU leaves (as a result of death or survival), one of the patients in the general ward is randomly chosen for admission to the newly vacated bed.

Greedy Policy (GP). This is an index policy, which gives priority according to the order determined by the differences $\phi_i^G - \phi_i$, where ϕ_i and ϕ_i^G denote the probability of death in the ICU and the general ward, respectively, for a patient in health stage i . Whenever an arriving patient finds the ICU full, the policy discharges the patient whose survival probability will have the smallest drop as a result of being treated in the general ward as opposed to the ICU. Similarly, when a patient leaves the ICU (as a result of death or survival), among the patients in the general ward, the patient with the most to benefit is chosen. Note that based on our mathematical analysis of the single-bed scenario, specifically Proposition 1, it might be reasonable to expect that GP would perform well when the patient demand is relatively low but when demand is high, as we assume in our simulation study, because the policy ignores the expected lengths of stay, we would not expect the policy to perform well.

Ratio Policy (RP). This is an index policy, which gives priority according to the order determined by $(\phi_i^G - \phi_i)/L_i$, where L_i is the expected length of stay for

patients in health stage i . Whenever an arriving patient finds the ICU full, the policy discharges one of the patients with the smallest expected drop in the survival probability divided by the expected length of ICU stay. Similarly, when a patient leaves the ICU (as a result of death or survival), among the patients in the general ward, the patient with the largest value of $(\phi_i^G - \phi_i)/L_i$ is chosen. Our mathematical analysis provides strong support for this heuristic, particularly when demand is high and thus one would expect good performance from this policy in the simulation study. Specifically, Proposition 1, which assumes the simplistic single-bed setting, finds that this policy is optimal when the arrival rate is sufficiently high. For the multibed scenario, we know from Theorem 2 that the optimal policy has a threshold structure, which would still be in line with RP, but we also have examples that show that RP is not optimal in general and that the optimal policy is state dependent. Nevertheless, Theorem 3 finds that under a condition for which RP and GP would be in complete agreement, RP would be optimal.

An important assumption underlying GP and RP is that we can observe each patient’s health stage precisely. Practically, however, this may not be possible. Patients could still be evolving according to some more sophisticated transition probability structure like the one we assumed in the simulation model, but we might only be able to do some rough classification and make decisions accordingly without knowing precisely in which health stage each patient is in. In fact, this would most likely be the case at least in the foreseeable future as it is very difficult if not impossible to come up with a classification system that perfectly captures patient evolution at a detailed level. To get a sense of how the policies we propose would perform in such a case, we also consider aggregated versions of GP and RP. They are aggregated in the sense that, as shown in Figure 3, if the patient is in one of the health stages 1, 2H, or 2L, the decision maker knows that the patient is in one of these stages but not the exact health stage. Thus, the decision maker assumes that the patient is in some aggregated stage A1. Similarly, if the patient is in one of the health

stages 3H, 3L, or 4, the decision maker, not knowing the exact health stage of the patient, assumes that the patient is in the aggregated stage A2. This means that the decision maker puts patients in only one of two health stages as in the case of our mathematical formulation. This allows us to consider a setting where the reality is complicated (as described in the simulation study) but the decision maker follows the policies suggested by our mathematical analysis. Note that using the aggregated stages requires estimation of transition probabilities among the aggregated stages A1 and A2 as well as the death and survival stages 0 and 5. We explain how the decision maker makes this estimation in Online Appendix A5.3.

Aggregated Greedy Policy (AGP). This policy is the same as GP except that it is applied over the aggregated classifications. When a patient from a particular aggregated stage is to be discharged or admitted from the general ward, one of the patients from that aggregated health stage is chosen randomly.


Aggregated Ratio Policy (ARP). This policy is the same as RP except that the policy is applied over the aggregated classifications. As in the case of AGP, when a patient from a particular aggregated stage is to be discharged or admitted from the general ward, one of the patients from that aggregated health stage is chosen randomly.

Aggregated Optimal Policy (AOP). When there are two health stages only and under the additional assumptions that when there are no patient readmissions from the general ward and the patient arrival process is stationary, we can determine the optimal policy by solving the MDP formulation described in Section 5. AOP basically uses the actions this optimal policy suggests (when the arrival rate is set to the current arrival rate in the simulation model). As in the cases of AGP and ARP, AOP randomly picks among the patients who belong to the same aggregated health stage.

6.3. Results of the Simulation Study

In the simulation study, we considered two DPCD values (3% and 5%), and three levels for the baseline ICU load (0.5, 0.8, and 1) as described in Section 6.1. Thus, in total, we considered six different combinations. The performance measure for each policy π (described in Section 6.2) was chosen to be the mortality rate, M_π , which we define to be the percentage of deaths among the patients who arrived at the ICU for possible admission during the 36-week period. We generated 30 transition probability scenarios for each one of the six DPCD-load pairs as described in

Figure 3. (Color online) Aggregated Two-Stage Transition Diagram for Patient Evolution in the ICU.



Section 6.1 and ran 100 replications for each scenario. In each replication, we randomly determined the initial state of the system. Specifically, the number of patients initially in the ICU was set assuming that the number is uniformly distributed over the integers from 0 to b and the health stage i of each patient is determined using the probability distribution $\{\theta_i, i \in \{1, 2L, 2H, 3L, 3H, 4\}\}$.

Using simulation results, we made pairwise performance comparisons between RP and every other policy. Specifically, we calculated the mean value for $M_\pi - M_{RP}$ for every policy π (over the 100 replications) for each scenario and constructed a 95% confidence interval for the mean difference. In Figures 4, 5, and 6, we provide these confidence intervals along with the box plots, where we also indicate the first and third quantiles, the minimum, and the maximum values.

From the figures, we can observe that RP has a superior performance overall when compared with the other policies. The good performance of RP is evident particularly when the comparison is made with respect to benchmark policies FCFS, RDP, GP, and AGP and the load on the ICU is very high. Given our mathematical analysis, the effect of system load on the performance of RP is not surprising. As we discussed earlier, when the system load is high, policies like GP, which exclusively takes into account the immediate benefit for the patients while ignoring system level factors such as the expected length of stay for the patients, are more likely to perform badly.

If we compare RP with the aggregated-type policies ARP and AOP, we observe that, even though these two policies perform better than the benchmarks, the performance of RP is again statistically better with the differences in the performances getting larger as the load on the system increases. Note that this comparison is important because as we discussed in Section 6.2, the model with which we are making decisions (e.g., our mathematical model) could be simpler than the reality (e.g., our simulation model as we assume in this paper). As the decision maker, we may not even know which specific health stage the patient is in but could only have some rough idea about the patient's health condition. With this comparison, we see that there is a benefit to knowing the health conditions of the patients in more detail, especially when the system is heavily loaded.

To get a better sense as to why RP performs well and its performance gets better with increased ICU load, recall Proposition 1 (particularly part (b)), which provides a necessary and sufficient condition for the optimality of RP for the special case where there is a single bed. According to the proposition, RP is optimal if the arrival probability exceeds a particular level, that is, if the inequality in Equation (5) is violated. This

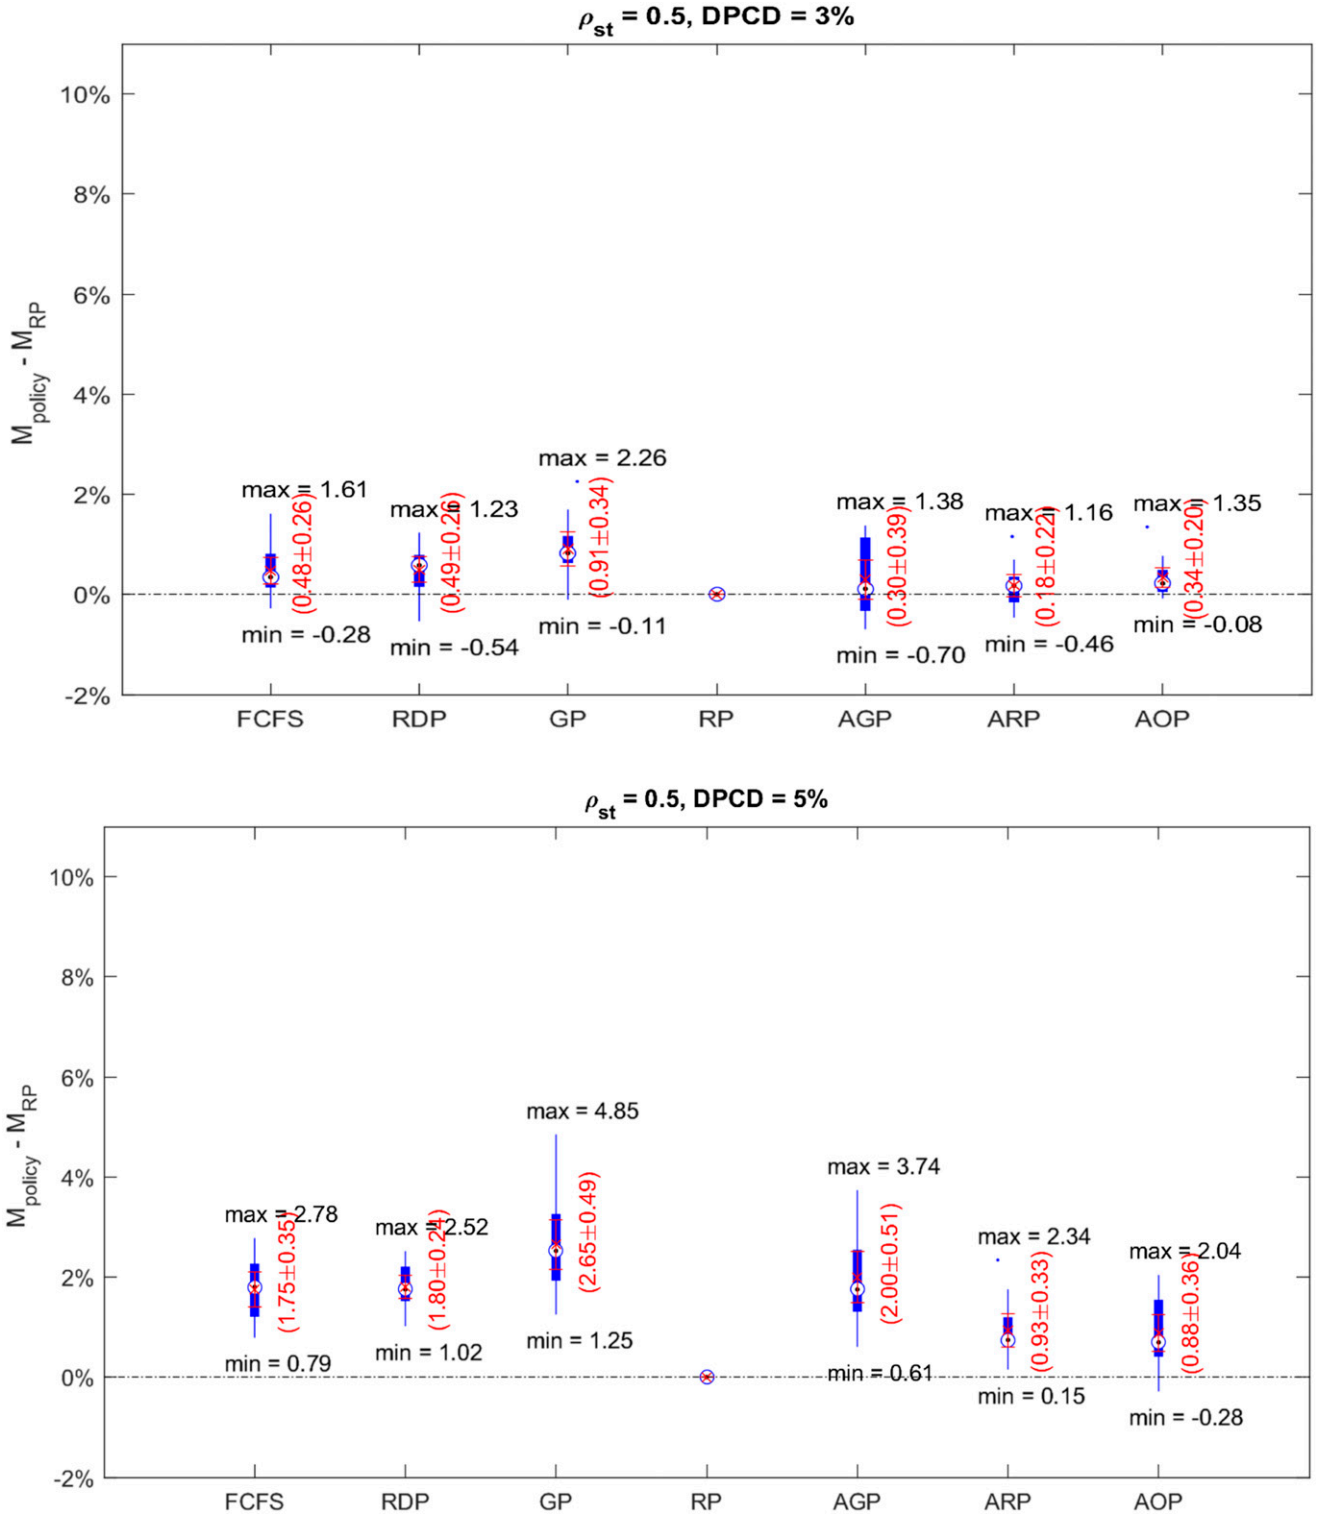
suggests that in general, RP could be more preferable when the ICU is highly loaded. In our simulation study, there are multiple beds, the arrival probability changes with time, and the inequality is written specifically for a model that assumes two health stages. Therefore, condition (5) is not well-defined in the context of our simulation model. However, one can still get some general sense for whether the arrival probability is high enough to favor RP by adjusting the arrival probability λ by λ/b , using the aggregated version of the transition probabilities under each scenario, and then determining the percentage of time the inequality in Equation (5) is violated. Following this procedure, we obtained Table EC.1 in Online Appendix A5.4.

As we can observe from the table, in many of the scenarios considered in our simulation study, the fraction of time the adjusted version of condition (5) is violated is either 1 or close to 1 providing some explanation as to why RP has such a good performance. It is also important to note that as the load on the system increases, the fractions under each scenario are also nondecreasing, which might explain why the performance of RP is more dominant when ICU load is higher.

Going back to Figures 4, 5, and 6, another observation we can make is that among the aggregated-type policies, ARP and AOP appear to perform better than AGP except for the case when the load on the ICU is the smallest. If we compare ARP with AOP, we see that even though the mean performance of AOP is better than that of ARP for all ICU load levels, the differences are not statistically significant. This suggests that even when patients can only be classified at the aggregate level as described earlier and thus RP is not an option, using the policy that is optimal (for our stylized formulation) may not be justified and the aggregated version of RP might be acceptable.

The observations that RP performs better than AOP and AOP does not seem to have a statistically strong advantage over ARP highlight two important questions that are closely intertwined: In searching for the discharge/admit policy to use in practice, can we restrict ourselves to state-independent policies? Given that the simulation study suggests that the best policy is state independent, does Theorem 2, which states that in general the optimal policy is of threshold type, have any practical value? For several reasons, it is difficult to provide definite answers to these questions. First of all, we know that AOP is optimal for our mathematical model but this obviously does not mean that it would continue to perform well when we change the underlying model from one with two health stages (as in the mathematical model) to one with six health stages and a more complex transition structure (as in the simulation model). It is not

Figure 4. (Color online) Pairwise Comparisons of the Differences in Average Mortality Rates Between Other Policies and Ratio Policy for Each Scenario with $\rho_{st} = 0.5$

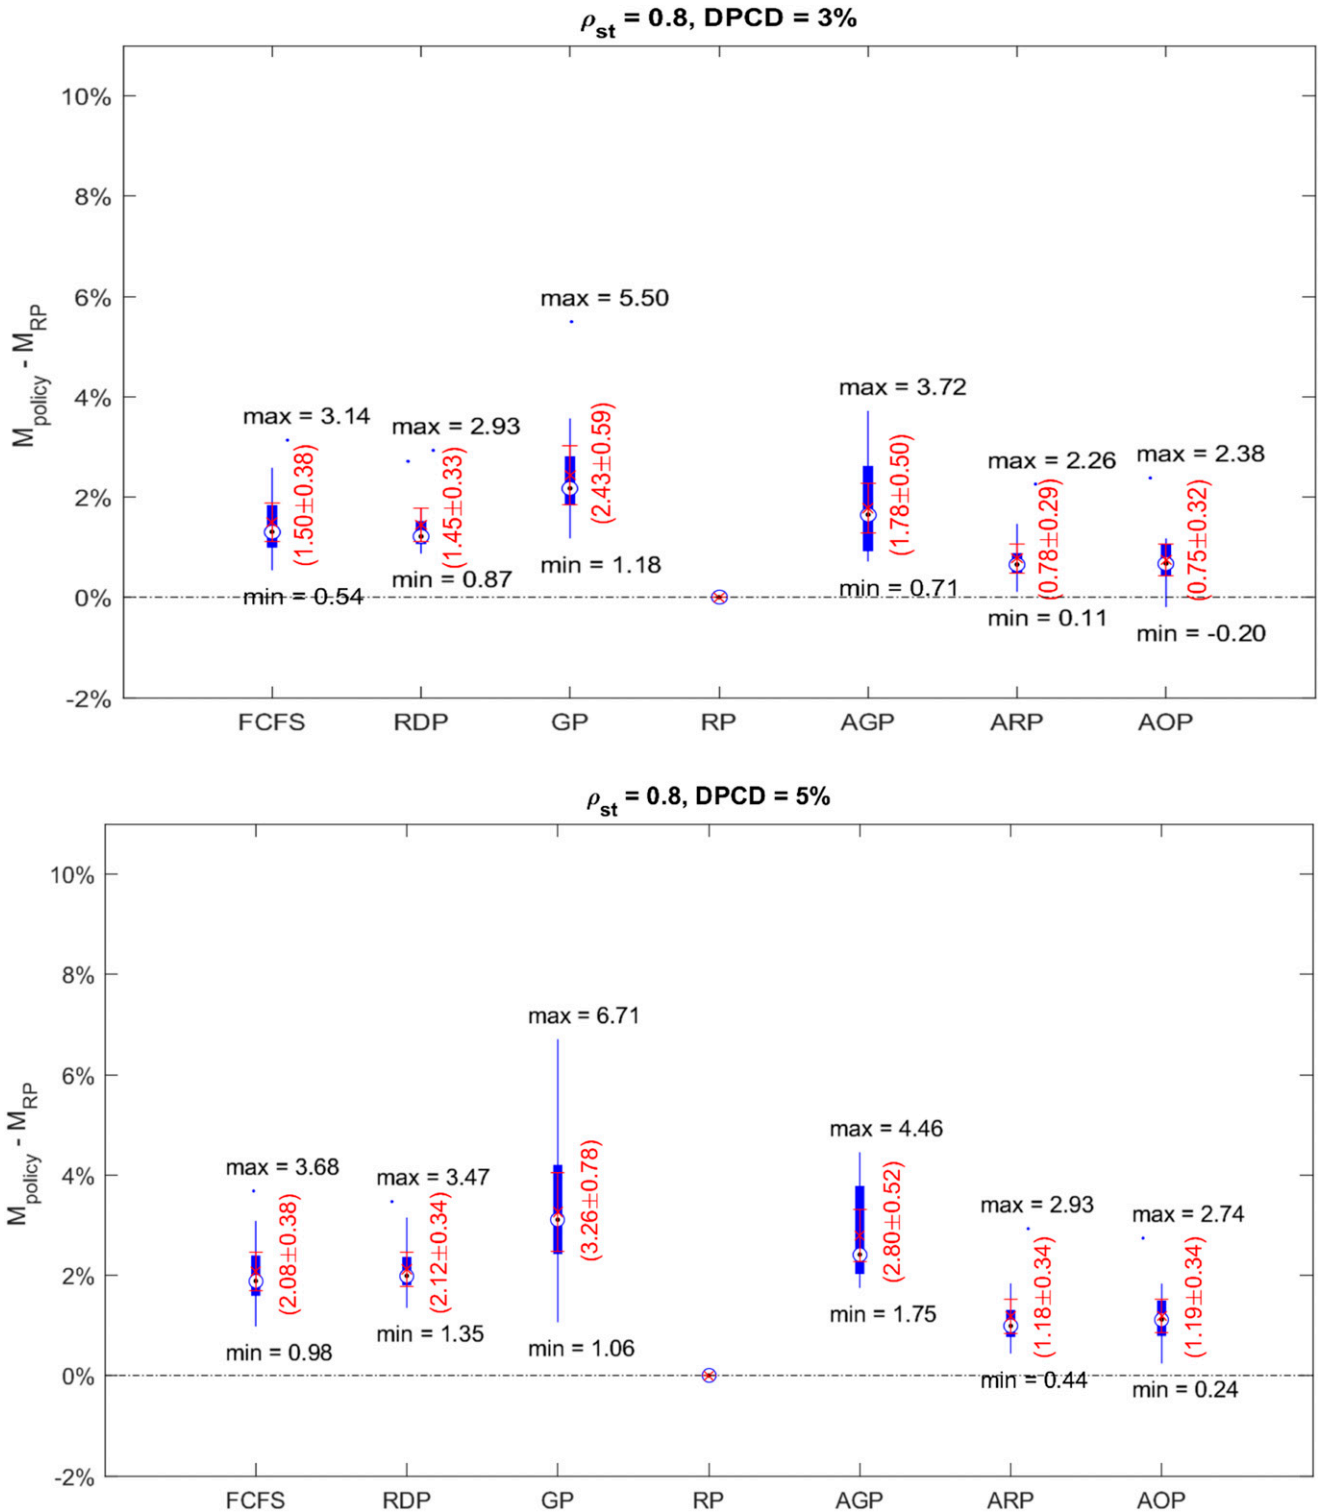


Note. Average is taken over 100 replications.

even clear whether AOP is the best among all the aggregated-type policies one could use for the model assumed in the simulation study. In short, AOP may not have performed as well as one would hope, but

this does not rule out the possibility of the existence of a different state-dependent policy that performs better. But more importantly, even though our patient health evolution formulation assumed in the

Figure 5. (Color online) Pairwise Comparisons of the Differences in Average Mortality Rates Between Other Policies and Ratio Policy for Each Scenario with $\rho_{st} = 0.8$




Note. Average is taken over 100 replications.

simulation study is highly likely to be an improvement over the one we assumed in the mathematical model, we do not know how well this particular model captures reality. As we discussed in detail in

Section 6.1, existing research on ICU patients is not at a level where we have a clear understanding of how ICU patients can be classified and how their health conditions evolve inside and outside the ICU. The

Figure 6. (Color online) Pairwise Comparisons of the Differences in Average Mortality Rates Between Other Policies and Ratio Policy for Each Scenario with $\rho_{st} = 1$



Note. Average is taken over 100 replications.

model we used in the simulation study is only one possibility among the many plausible. Therefore, even though our simulation study provides some

useful insights and directions for future work, it would not be reasonable to make immediate generalizations from our observations. With more research

in this area, we will have an increasingly better understanding of ICU patients and be able to develop models that are increasingly better representations of reality. It is possible that with changes in the health evolution model, performances of the policies relative to each other will also change and it will be prudent to construct potentially good new state-dependent policies and investigate their performances. The insights that come out of the optimal policy characterizations given in Theorem 2, which describe how the composition of the patients in the ICU should influence admit/discharge decisions, can be helpful in the construction of such policies.

Even though our simulation study cannot provide a definite answer to the question of which policy would work better in practice, the fact that RP had the best performance is good news. The policy is simple, easily generalizable, intuitive, and does not need to keep track of system state information. It is also important to note that the policy only requires the estimation of expected net benefits and the expected lengths of stay for each health stage, not the individual transition probabilities. This not only makes it much easier to implement RP in practice, but also means that the policy is highly robust to transition probability estimates and the assumptions made regarding the underlying patient health and transition formulation.

Finally, in this section, we investigate patients' lengths of stay in the ICU under each policy. Figures A2, A3, and A4 given in Online Appendix A5.5 summarize the results of our analysis. We can observe from the figures that if we leave aside FCFS, there are no notable differences between the policies with respect to average lengths of ICU stay. Long lengths of stay under FCFS is not surprising because under that policy patients leave the ICU only when they are dead or they reach the survival stage. They are never discharged early to accommodate other patients. Under any of the other policies, patients can be discharged from the ICU even though they still need ICU care and this results in shorter lengths of stay. We can also observe from the figures that the lengths of stay under every policy except FCFS decrease as the ICU load increases. This is because except in the case of FCFS, the more patients there are in need of ICU, the higher the chances that any given patient's ICU stay is cut short, which ultimately leads to shorter average lengths of stay.

7. Conclusions

Many studies reported that the number of ICU beds in many parts of the United States and the rest of the world are in short supply to sufficiently meet the daily ICU demand. It is frequently the case that a patient who is in a relatively less critical condition is discharged early to make room for another patient who is deemed more critical. Although this bed shortage

problem arises even under daily operating conditions, it is natural to expect the problem to get worse in case of an event like an influenza epidemic, which causes a significantly increased number of patients in need of an ICU bed. It is thus highly important to investigate how ICU capacity can be managed efficiently by allocating the available beds to the patients in a way that the greatest good is achieved for the greatest number of the patients. Our goal in this paper has been to provide insights into and develop policies for making such allocation decisions.

What mainly sets our analysis apart from prior work is that in our model we allow the patients to move from one health stage to another and allocation decisions are made based on the patients' updated health conditions. This formulation captures an important feature of the actual problem and nicely fits with the triage protocol proposed by Christian et al. (2006). But more importantly, the model allowed us to go deeper and establish properties that appear to be difficult to identify using formulations considered in prior work. For example, we were able to provide analytical results for the case where patients who have higher expected ICU benefits also have longer expected length of stay.

Our analysis of the single-bed scenario led to interesting insights into how optimal decisions depend on the patients' expected ICU benefit, expected length of stay, and the patient load on the system. We found that when patients who are expected to benefit more from ICU treatment also have longer expected lengths of stay, those patients should get higher priority only if the overall patient demand is below a certain level. This is because when beds are in high demand, prioritizing those patients (who are expected to occupy the beds longer) would require turning too many patients away from the ICU that it becomes more preferable to adopt a policy that has quicker bed turnaround times even though the expected net benefit is smaller for every admitted patient. More generally, when the ICU has finitely many beds, we found that the optimal policy aims for an ideal mix in the ICU so as to hit the right balance between the overall expected net ICU benefit per patient and length of stay. That is, in general, the optimal policy for prioritizing among patients depends on the mix of patients in the ICU.

Considering the complexity of the actual decision problem we are interested in, the mathematical model we analyze in this paper is stylized and therefore it is natural to question the generalizability of the main insights. Indeed, our simulation study, which, unlike the mathematical model, allows readmissions from the general ward and considers a more complex patient health evolution formulation, suggests that there may not be a justification for searching for a complex policy that prioritizes based on the patient mix in the ICU.

On the other hand, the simulation study also shows that some of the policies that are proposed based on our mathematical analysis performs well even under the more general conditions of the simulation model. The fact is that not enough is known about the ICU patients for us to be able to construct a very realistic description of patient evolution. The more complex model considered in the simulation study is another simplification at best. Therefore, not only the conjectures on the generalizability of the policies should be taken with a grain of salt, but one should also not quickly conclude from our simulation study that in practice there is no need to consider policies that take patient mix into account. Nevertheless, our results provide some confidence that despite the complexity of the decision problem in practice, relatively simple policies might work well, and the paper provides some useful guidance for what future research and data collection efforts should focus on in order to develop useful patient classification and triage protocols and ultimately decision support tools that can be implemented in practice.

Acknowledgments

The authors thank the department editor, the associate editor, and the referees whose comments and suggestions greatly improved this paper.

References

- Armony M, Chan CW, Zhu B (2018) Critical care capacity management: Understanding the role of a step down unit. *Production Oper. Management* 27(5):859–883.
- Cardoso LT, Grion CM, Matsuo T, Anami EH, Kauss IA, Seko L, Bonametti AM (2011) Impact of delayed admission to intensive care units on mortality of critically ill patients: A cohort study. *Critical Care* 15(1):Article R28.
- Chalfin DB, Trzeciak S, Likourezos A, Baumann BM, Dellinger RP, DELAY-ED Study Group (2007) Impact of delayed transfer of critically ill patients from the emergency department to the intensive care unit. *Critical Care Medicine* 35(6):1477–1483.
- Chan CW, Yom-Tov G, Escobar GJ (2014) When to use speedup: An examination of service systems with returns. *Oper. Res.* 62(2):462–482.
- Chan CW, Farias VF, Bambos N, Escobar GJ (2012) Optimizing intensive care unit discharge decisions with patient readmissions. *Oper. Res.* 60(6):1323–1341.
- Christian MD, Hawryluck L, Wax RS, Cook T, Lazar NM, Herridge MS, Muller MP, Gowans DR, Fortier W, Burkle FM (2006) Development of a triage protocol for critical care during an influenza pandemic. *Canadian Medical Assoc. J.* 175(11):1377–1381.
- Dobson G, Lee HH, Pinker E (2010) A model of ICU bumping. *Oper. Res.* 58(6):1564–1576.
- Ghassemi M, Naumann T, Doshi-Velez F, Brimmer N, Joshi R, Rumshisky A, Szolovits P (2014) Unfolding physiological state: Mortality modelling in intensive care units. *Proc. 20th ACM SIGKDD Internat. Conf. Knowledge Discovery and Data Mining (ACM, New York)*, 75–84.
- Gortzis LG, Sakellaropoulos F, Ilias I, Stamoulis K, Dimopoulou I (2008) Predicting ICU survival: a meta-level approach. *BMC Health Service Res.* 8(1):Article 157.
- Kc DS, Terwiesch C (2012) An econometric analysis of patient flows in the cardiac intensive care unit. *Manufacturing Service Oper. Management* 14(1):50–65.
- Kim SH, Chan CW, Olivares M, Escobar G (2014) ICU admission control: An empirical study of capacity allocation and its implication for patient outcomes. *Management Sci.* 61(1):19–38.
- Kumar A, Zarychanski R, Pinto R, Cook DJ, Marshall J, Lacroix J, Stelfox T, Bagshaw S, Choong K, Lamontagne F (2009) Critically ill patients with 2009 influenza A (H1N1) infection in Canada. *J. Amer. Medical Assoc.* 302(17):1872–1879.
- 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(10):1345–1355.
- Örmeci EL, Burnetas A (2005) Dynamic admission control for loss systems with batch arrivals. *Adv. Appl. Probab.* 37(4):915–937.
- Örmeci EL, Burnetas A, van der Wal J (2001) Admission policies for a two class loss system. *Stochastic Models* 17(4):513–540.
- Puterman ML (2005) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley & Sons, Hoboken, NJ).
- Rocker G, Cook D, Sjøkvist P, Weaver B, Finfer S, McDonald E, Marshall J, et al. (2004) Clinician predictions of intensive care unit mortality. *Critical Care Medicine* 32(5):1149–1154.
- Shmueli A, Sprung CL (2005) Assessing the in-hospital survival benefits of intensive care. *Internat. J. Tech. Assessment Health Care* 21(1):66–72.
- Shmueli A, Sprung CL, Kaplan EH (2003) Optimizing admissions to an intensive care unit. *Health Care Management Sci.* 6(3):131–136.
- Sinuff T, Kahnemouli K, Cook DJ, Luce JM, Levy MM (2004) Rationing critical care beds: A systematic review. *Critical Care Medicine* 32(7):1588–1597.
- Strand K, Flaatten H (2008) Severity scoring in the ICU: A review. *Acta Anaesthesiologica Scandinavica* 52(4):467–478.
- Ulukus MY, Güllü R, Örmeci EL (2011) Admission and termination control of a two class loss system. *Stochastic Models* 27(1):2–25.
- Vincent JL, Moreno R, Takala J, Willatts S, De Mendonça A, Bruining H, Reinhart CK, Suter P, Thijs LG (1996) The SOFA (sepsis-related organ failure assessment) score to describe organ dysfunction/failure. *Intensive Care Medicine* 22(7):707–710.
- Wagner J, Gabler NB, Ratcliffe SJ, Brown SE, Strom BL, Halpern SD (2013) Outcomes among patients discharged from busy intensive care units. *Ann. Internal Medicine* 159(7):447–455.
- 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(5):1297–1310.

Huiyin Ouyang is an assistant professor of innovation and information management in the Faculty of Business and Economics at the University of Hong Kong. Her research interests are in stochastic modeling and healthcare operations, with a focus on resource allocation and priority assignment in healthcare services.

Nilay Tanık Argon is a professor of statistics and operations research at the University of North Carolina. Her research interests are in stochastic modeling of manufacturing and service systems, queueing systems, healthcare operations, and statistical output analysis for computer simulation.

Serhan Ziya is a professor of statistics and operations research at the University of North Carolina. His research interests are in service operations, with a focus on healthcare operations, queueing systems, revenue management, and pricing under inventory considerations.