

Elective Patient Admission and Scheduling under Multiple Resource Constraints

Christiane Barz, Kumar Rajaram

Anderson School of Management, University of California, Los Angeles, CA 90095,
christiane.barz@anderson.ucla.edu, krajaram@anderson.ucla.edu

We consider a patient admission problem to a hospital with multiple resource constraints (e.g. OR and beds) and a stochastic evolution of patient care requirements across multiple resources. There is a small but significant proportion of emergency patients who arrive randomly and have to be accepted at the hospital. However, the hospital needs to decide whether to accept, postpone or even reject the admission from a random stream of non-emergency elective patients. We formulate the control process as a Markov Decision Process to maximize expected contribution net of overbooking costs, develop bounds using approximate dynamic programming and use them to construct heuristics. We test our methods on data from the Ronald Reagan UCLA Medical Center and find that our intuitive newsvendor-based heuristic performs well across all scenarios.

Key words: Patient admission, patient scheduling, multiple resources, Markov decision process, Approximate dynamic programming

History: Received: September 2014, accepted: April 2015 by Christopher S. Tang after one revision

1. Introduction

The health care industry is a major component of the US Economy, representing around 17% the GDP (OECD 2012). The US spends more on healthcare as a portion of its GDP and on a per capita basis (\$7146) than any other nation in the world (WHO 2011). In addition, health care costs are rising exponentially (SSAB 2009). Public funds cover about 45% of these costs and the amount of public funding in healthcare is expected to double by 2050 (Gupta and Denton 2008). This cost explosion in health care over the last decades has spurred the use of operations management techniques for various problems in this field ranging from process design, capacity allocation, aggregate planning, admissions control and appointment scheduling. Pierskalla and Brailer (1994) provide an excellent overview of these problems.

Of particular importance is the patient admission control problem as this has implications on virtually every other problem faced in a health care setting, such as clinics or hospitals. This problem is concerned with deciding which patient to admit to the hospital and at what time. This problem is important from a patient perspective as it defines the quality, access and time

component of the service. It is important from a hospital perspective as different patients bring different revenues and cause different costs. The patient mix that arises from an admission policy hence defines the revenues, costs and ultimately profitability of the entire hospital. The solution to the admission control problem is particularly complicated due the following five reasons. First, many hospitals face multiple resource constraints and depending upon the patient mix any one of them can become the bottleneck (Duda et al. 2013). As a consequence, the analysis cannot be restricted to the consideration of one resource, but the resource consumption of all potential bottlenecks must be modeled. Second, the arrivals of patients to the hospital are not predictable or deterministic. Third, only the admission of some patients can be controlled, while other patients must always be accepted immediately. We will refer to the latter patients, the uncontrolled stream of admissions, as emergency patients although it may comprise much more than just life and death emergencies such as those mandated by insurance companies, government regulations, or cases sought out for teaching or research purposes. Admission and scheduling decisions can only be made for elective patients. Fourth, resource usage at each stage of the hospital varies across patients even for the same procedure and the expected future resource usage needs to be updated over time. Finally, current admissions constrain the admissions of future patients to the hospital.

Our goal is to assist decision makers with patient admission and scheduling, which we view as an input to other scheduling decisions throughout the hospital. In this sense, our patient scheduling model accounts for overall capacities of all resources but does not incorporate or replace daily OR scheduling or bed assignment decisions. The patient mix resulting from our model should be viewed as an input to such models. A result of our model formulation are estimates of the opportunity costs of units of all potential bottlenecks for given demand from elective patients. These values can be used to adjust the contribution from different patient types and hence obtain a ranking or priority rules for elective admissions. Such a ranking can then be used to decide how to expand or contract various practices and to decide block time allocations. The opportunity costs can also be used to decide how many units of each resource should be blocked for emergency patients.

We formulate the elective patient admission control problem as a Markov Decision Process (MDP). This approach is useful to model sequential decision problems with stochastic characteristics, which possess the Markovian property (i.e. future states and decisions are independent from past states, given the knowledge of the present state of the system) and there is a possibility of observing the system state at decision instants equally spaced over time (such as a day). This seems particularly relevant in our context as how a patient responds to a treatment and the corresponding care requirements are stochastic, while doctors evaluate and make treatment decisions on patients each day based on the current health state of the patient (Schaefer et al. 2004).

Modeling a Markov decision process that incorporates multiple resource constraints and a stochastic evolution of patients' care requirements leads to an MDP formulation that has a state space that is too large to allow for a direct solution. To overcome this, we resort to Approximate Dynamic Programming (ADP) to develop heuristics for this problem and to construct a bound to evaluate the quality of the heuristics. We test our heuristics with real data from the Ronald Reagan University of California at Los Angeles (RRUCLA) Medical Center and find that our intuitive newsvendor heuristic significantly outperforms the current practitioner based heuristics. This performance improves further in resource constrained environments with increasing demand from elective patients.

1.1. Related Literature

Our paper can be related to several streams of literature in the area of application and in the methodologies employed. In terms of the application, our work is connected to work in health care management on appointment scheduling/patient admission and treatment planning, as well as revenue management for hospitals and hotels.

The basic problem of patient admission and appointment scheduling is concerned with assigning appointment times to patient requests such that patient waiting time, server idle time, and/or server overtime is minimized. Usually, it is assumed that patients only need access to one (type of) resource, such as the OR, a diagnostic facility, or a bed, at the time of their appointment. An early discussion of when to schedule elective patients in the presence of mandatory (or emergency) patients, which must be accepted immediately, when the number of beds is limited is given in Kolesar (1970). Overviews of appointment scheduling research can be found in Cayirli and Veral (2003), Mondschein and Weintraub (2003) and Gupta and Denton (2008). More recent works include Patrick et al. (2008), Liu et al. (2009), Robinson and Chen (2010), Gocgun et al. (2011), Saure et al. (2012) and Patrick (2012).

Most literature in patient admission and scheduling has assumed that the hospital has only one potential bottleneck, see e.g. Gerchak et al. (1996), Shmueli et al. (2003), Helm et al. (2011), or Kim et al. (2013). As the process analysis by Duda et al. (2013) shows, the hospital we considered is constrained by more than one resource. Given a certain patient mix, there is usually only one bottleneck. For different patient mixes, however, different resources constrain the capacity of the facility leading to different bottlenecks. Since in patient admission, the patient mix is determined dynamically, it is critical that all potential bottlenecks are considered. Acknowledging that there can be more than one constraining resource in a hospital, Kusters and Groot (1996) present a statistical model for the prediction of resource availability. Patient admission problems that consider more than just one constraining resource were suggested by Adan and Vissers (2002) and Vissers et

al. (2005). While they view the problem as deterministic and model it as a mixed integer program, simulation is used to find good admission policies in Oddoye et al. (2009). Adan et al. (2011) formulate a two-stage planning procedure for a hospital with four resources and stochastic length of stay to minimize the deviations of the resource consumption from a given target utilization. Within this stream of literature, the papers of Helm and Van Oyen (2010) and Huh et al. (2013) are most closely related to ours. Focussing on the tradeoff between throughput and blockage, Helm and Van Oyen (2010) restrict their attention to beds in different wards and do not model other resources. To model demand for beds in different wards, they consider an open network of infinite server queues; admission decisions are obtained by a mixed integer program that aims at a tradeoff between throughput and blockage. They test their approach with real data but cannot determine how close their solution is to the true optimum. In contrast to focussing on throughputs, we maximize contribution from patients using an approximate dynamic programming framework. This approach is flexible enough to allow for any set of constraining resources, inspires an intuitive heuristic and naturally provides an upper bound to our objective function. Huh et al. (2013) recently suggested a Markov decision process formulation for a multiresource allocation problem with two job types (elective and emergency). To minimize the discounted sum of waiting costs and lost revenue due to renegeing elective patients over a finite horizon, they allow for evolving and possibly correlated information about the distribution of demand and capacity. In contrast to our paper, they do not allow for patient heterogeneity in resource consumption within the group of emergency and elective patients and assume that patients only consume resources for one period. Further, their model specifies how many elective patients to admit for the current period but not the timing or sequence of service.

There has been a significant body of literature on treatment planning. For an overview on treatment planning, see Sox et al. (1988). While these earlier models mostly focused on decision trees, more recent approaches favor the use of Markov decision processes, as in Sonnenberg and Beck (1993), Naimark et al. (1997), Schaefer et al. (2004), and Alagoz et al. (2010). In our work, we do not model treatment planning while making admissions decisions as they are typically made by different entities at different points of time. For example, the hospital administration makes decisions on admissions and scheduling, while treatment plans are exclusively developed by doctors after assessing the patient once they have been admitted. We assume that for each patient, a finite stage Markov decision process can be formulated to find the optimal treatment plan. Given the optimal policy for treatment, the patient's recovery hence follows a Markov chain. Kapadia et al. (1985) suggest how to estimate the transition probabilities in such a context. To our knowledge, the only other paper in admission planning and scheduling that models patients in terms of their recovery state is Nadal Nunes et al. (2009). Trying to achieve a target utilization, they, however, do

not allow for emergency patients or scheduling elective patients into future time periods. Although they allow for stochastic resource usage, their admission model assumes that demand from elective patients is constant and deterministic in each period. From a practical point of view, the complexity of their model does not lead to any managerial insights and can only be applied to extremely small problem instances due to the curse of dimensionality.

Another related application is the field of revenue management. Capacity control in revenue management deals with the control of the selling process of a perishable resource. Since hospital resources are services that cannot be stored for future time periods, we can view patients as customers that ask for a certain combination of perishable resources at a certain price. However, in contrast to classical network revenue management, we have the following important distinctions: 1) uncontrolled demand in the form of emergency patients, 2) uncertain resource requirements because initially it is uncertain how much of each resource a patient needs each day, 3) flexibility in the assignment of resources since, within a certain planning horizon, we can decide on what day a patient should be admitted, and 4) an infinite planning horizon, where resources perish sequentially. For overviews on revenue management, and in particular network revenue management, see Talluri and van Ryzin (2004) and Chiang et al. (2007). Hotel revenue management can be viewed as a special kind of network problem that is close in terms of application. An overview is presented in Pullman and Rodgers (2010). Early work includes Bitran and Mondschein (1995) and Bitran and Gilbert (1996). Recently, a hospital revenue management problem was suggested by Ayvaz and Huh (2010). While they also model emergency and elective patients, their understanding of the two groups differs from ours since emergency patients can be rejected if the hospital runs out of capacity and elective patients need not be scheduled for a future time period but wait until there is capacity to serve them. Further, they assume that all patients need exactly one unit of one single resource at the day of admission (and no resources on later days). Patients hence have deterministic and equal resource usage.

In terms of methodology, overviews on approximate dynamic programming can be found in Bertsekas and Tsitsiklis (1999), Bertsekas (2005), and Powell (2007). These overviews cover mainly simulation-based approaches to approximate dynamic programming. However, we use the linear programming (LP) approach to approximate dynamic programming. The LP approach was first suggested by Schweitzer and Seidmann (1985), but has gained attention only in the past decade. Recent papers include de Farias and Van Roy (2003), Adelman (2003), Adelman (2004), de Farias and Van Roy (2004), de Farias and Van Roy (2006), Adelman (2007), and Adelman and Mersereau (2008). The LP approach has been applied in a patient admission/scheduling problem to a diagnostic resource by Patrick et al. (2008) and for chemotherapy in Gocgun and Puterman (2014). In

contrast to our work, Patrick et al. (2008) aims at controlling waiting times (rather than contribution) by rejecting or postponing the treatment of outpatients. Patients differ in terms of priority but all require exactly one unit of one resource unlike our model that can handle multiple and time varying resource requirements. Similarly, patients in Gocgun and Puterman (2014) all require one time slot and purely aim at scheduling patients within a given target window.

1.2. Contributions

This paper makes the following contributions to both the application and methodology domains:

1. To the best of our knowledge this paper enhances existing patient admission models by considering the expected net contribution maximizing elective patient admission and scheduling problem under stochastic evolution of patients health and care requirement with multiple resource constraints or bottlenecks which change depending upon elective admissions or patient mix. As stated in Gupta and Denton (2008), both these aspects are important and open challenges in the literature. In this context, we provide a novel formulation of an average contribution maximizing Markov decision process for this problem which schedules elective admissions in current and future time periods.

2. We use concepts from ADP to address the curse of dimensionality inherent in this realistic but more complicated formulation. In particular, we use an affine approximation of the bias function to obtain approximate values for the marginal cost of using one unit of each resource the hospital provides to its patients. Since the resulting linear problem is still hard to solve, we show how a judiciously chosen state space extension can lead to a more tractable upper bound without changing the approximation architecture. This formulation determines approximate values for the marginal cost of using one unit of each resource the hospital provides to its patients, connecting our ADP approximation to a newsvendor model.

3. We use the established connection to the newsvendor model to suggest a novel heuristic, which we compare to standard price-directed heuristics in a numerical study. The heuristic uses some key insights gained through the approximation and has the advantage of being less computationally intensive than the price-directed heuristics. Further, it is easy to communicate to practitioners as it can be viewed as an extension of a rule that is already used in practice. These heuristics also provide lower bounds for this problem. We apply our methods to real data from the RRUCLA medical center and show that the performance of the newsvendor-inspired heuristic is very competitive, outperforming greedy strategies and policies that were reported to us from practice by far.

In the next section, we introduce the model and formulate the optimization problem as a Markov decision process. In Section 3, we suggest an intuitive upper bound problem based on the deterministic version of this problem. In Section 4, we formulate the optimality equations, introduce

the ADP approximation, prove structural results, and suggest an algorithm to solve for the optimal approximation parameters. This provides an improved upper bound. In Section 5, we suggest heuristics based on these approximation parameters. In Section 6, we analyze a small example in detail and also demonstrate the performance of our methods using data from the RRUCLA Medical Center. In Section 7, we offer conclusions and provide future research directions.

2. Model Formulation

In this section, we formulate the elective patient admission model. Throughout, we will denote the set of natural numbers including 0 by \mathbb{N}_0 and the set of real numbers by \mathbb{R} . The hospital provides service in the form of R different resources. Each day, there is a capacity $c_r \in \mathbb{N}_0$ of resource $r = 1, \dots, R$. Capacity that is not used that day cannot be stored for future use but perishes.

We will introduce and model the dynamically changing health state of individual patients in the hospital granularly to make the model inputs intuitive to practitioners but we make a few simplifying modeling assumptions in the admission process. First, we assume that hospital utilization has no impact on the resource usage of the patients. Specifically, we assume that once the patient is admitted, the hospital will provide the best service possible, and utilization or newly incoming requests have no impact on the treatment of currently admitted patients. In practice, it may be possible to release patients earlier than recommended if hospital beds are scarce (see e.g. Diwas and Terwiesch (2009) and Diwas and Terwiesch (2012)). However, consistent with medical ethics, we do not model such quality-access-tradeoffs in patient treatment but view those effects as a result of tactical countermeasures. Because of this assumption, it might be unavoidable that more units of a resource are requested on a given day than what is available. We will allow for these situations based on practical considerations, but they lead to overtime costs (if, for instance, more OR time is used than planned), the loss of patient goodwill (if the lack of a regular bed requires the patient to sleep on a hospital gurney in a hallway), and the cost of inferior care (if a patient is assigned to a bed in a wing of a different specialty). All these costs are captured by the appropriate choice of the penalty cost $0 \leq \pi_r < \infty$ per unit of resource r per day which we use to discourage this overuse of resources. Second, we assume that emergency patients must always be accepted. Although ambulance diversion might reduce the stream of newly arriving emergency patients and is reported to be prevalent in practice (see e.g. Allon et al. (2013)), it is generally viewed as an unpopular measure of last resort. Further, it can only reduce not fully eliminate new arrivals since a hospital must always treat walk-in patients to the emergency department. We hence do not explicitly model this tactical countermeasure. Ambulance diversion can, however, be viewed as one way of reducing overtime. The cost of diverting an emergency patient to a different hospital is then accounted for by the corresponding penalty costs. As a result of these two assumptions, we allow for unrestricted overbooking in the analysis below.

2.1. The Patients

Each day, patients arrive in two different forms: emergency and elective patients. Emergency patients are always admitted immediately. For elective patients, we allow the hospital administration to determine admission and scheduling. Depending on the type of elective patient, there might be a certain time frame (such as “in the next two weeks”) during which a patient may be admitted. At the time the patient or their referring doctor asks for admission, they must be told when to come to the hospital for admission, or the patient must be referred to another hospital. We will call the latter decision a rejection.

Patients admitted to the hospital use resources in a particular way over time. At the time of admission, however, the exact resource requirements over time might not be known because only a limited amount of knowledge about the patient’s health condition is available. We call this knowledge the admission diagnosis j , $j = 1, \dots, J$. Patients with the same admission diagnosis might still react differently to the same treatment during the early days of their stay and complications that are observed with one patient need not be observed with the other. We therefore do not assume that all patients with a given admission diagnosis have the same resource requirements over time but allow them to differ for different patients (as indicated by the literature on medical decision making e.g. by Schaefer et al. (2004)).

In the following, we explain how individual patients are modeled. Although the exact resource requirements might not be known at the time of admission, we assume that a distribution of resource requirements is available for each patient. In particular, we assume that the requirements of any individual patient with admission diagnosis j on day n after admission can be described by a Markov chain $\{Z^n, n \in \mathbb{N}_0\}$ with initial state 0_j , and a state space that is composed of one absorbing state \diamond and a finite number of transient states $\{0_j, 1_j, \dots, M_j\}$. Transient states represent health states that determine the resource requirements of a patient, the absorbing state represents their discharge from the hospital. We do not differentiate between different modes of discharge (home, rehabilitation, transfer, or death) since all that matters to our problem is that no further resources are required for this patient. Transition probabilities are given as $p_{z^n z^{n+1}} = P(Z^{n+1} = z^{n+1} | Z^n = z^n)$. As \diamond is absorbing, we have $p_{\diamond\diamond} = 1$.

The function $u_r(z)$ assigns resource requirements to health states z . A patient in state z requires $u_r(z) \in \mathbb{N}_0$ units of resource r on the current day. We expect that a patient with initial diagnosis j will need $E[u_r(Z^n) | Z^0 = 0_j]$ units of the r th resource n days after admission. For each patient, we assume that we can observe their current state so that the resource requirements on the current day are known but the requirements of future days are only known in distribution. We denote by $n = 0$ the day of admission. Since state \diamond represents a discharged patient that no longer needs resources provided by the hospital, we have $u_r(\diamond) = 0$ for all $r = 1, \dots, R$. We assume that there is

a maximum number of days N a patient might stay in the hospital, i.e. there exists an $N \in \mathbb{N}$ with $P(Z^N = \diamond | Z^0 = 0_j) = 1$ for all admission diagnosis $j = 1, \dots, J$.

The above assumptions specify the Markov chain used to model the stay of the patients in the hospital. These assumptions are quite consistent with actual practice. To illustrate, consider a patient of type j represented in the box of the Markov chain in Figure 1. In this example, consider the two resources: OR time ($r = 1$) and surgical beds ($r = 2$), i.e. $R = 2$. A patient of type j requires 5 units of resource 1 and 1 unit of resource 2 on the day of admission. There is a 0.90 probability that no complication occurs. In that case, the patient needs to be monitored for 2 days after the original surgery. There is, however, a 0.01 probability that the patient will not survive the surgery and will not require further resources. Finally, there is a probability of 0.09 that they survive but suffer from a complication, which results in a follow up surgery using 2 units in the OR on the next day. In case of a complication, there is a 0.05 probability that the patient will not survive the second surgery. If the patient survives, they need a bed for another 2 days after the second surgery.

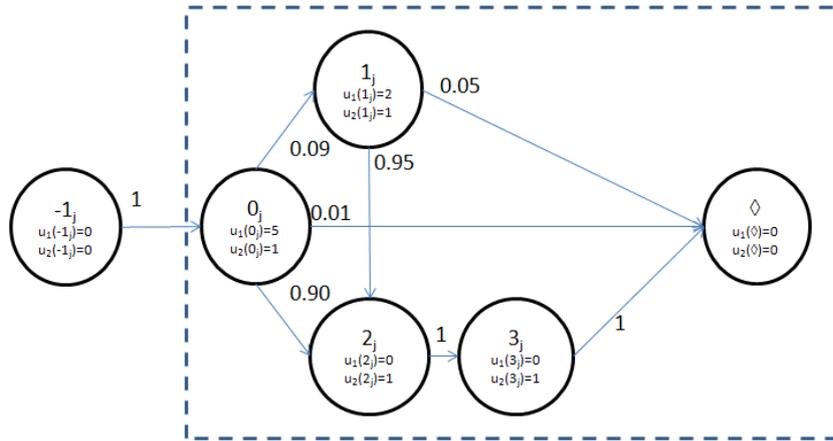


Figure 1 Example of a patient representation.

We assume that there only is a finite number of admission diagnosis, $j = 1, \dots, J$. Elective patients with the same admission diagnosis can bring different expected contributions f to the hospital depending on their insurance, even if they use the exact same resources. Further, they can be scheduled any time within a certain time horizon, i.e. within the next t days. This horizon may depend on the admission diagnosis and possibly other factors. We denote the combination of initial patient state as given by the admission diagnosis j , expected contribution f , and horizon t as the type (j_i, f_i, t_i) of an elective patient and assume that there is only a finite number of types, $i = 1, \dots, I$. Although emergency patients also bring some contribution to the hospital, we do not explicitly model their contribution since this part of the process cannot be controlled.

2.2. Scheduling Elective Patients

Each day, the hospital has to decide about the admission and scheduling of elective patients before emergency patients are observed. To represent a patient of type (j_i, f_i, t_i) that has been scheduled for admission in $\tau \leq t$ time periods, we introduce patient states $-t_i, \dots, -1_i$. A patient scheduled for admission in $\tau \leq t_i$ time periods diagnosed with j is in state $-\tau_i$. Although our model is flexible enough to allow for cancelations by elective patients in general, for the sake of simplicity we restrict ourselves to no cancelations in the following. As a consequence, we have $p_{(-\tau_i)(-\tau+1)_i} = 1$. Consequently, the full state space of an elective patient of type i with diagnosis j is $\{-t_i, \dots, -1_i, 0_{j_i}, \dots, M_{j_i}, \diamond\}$. Figure 1 represents a patient type with $t = 1$. Emergency patients with admission diagnosis j only have the subset $\{0_j, \dots, M_j, \diamond\}$ as their state space. Let $\mathcal{Z}' = \bigcup_{i=1}^I \{-t_i, \dots, -1_i, 0_{j_i}, \dots, M_{j_i}, \diamond\} \cup \bigcup_{j=1}^J \{0_j, \dots, M_j, \diamond\}$ be the union of all patient state spaces.

We assume the daily arrival distribution is known. Specifically, letting D_i be the random number of elective patients of type i asking for admission, and X_j the number of emergency patients with admission diagnosis j , we know $P(\vec{D} = \vec{d})$ and $P(\vec{X} = \vec{x})$ for random $\vec{D} = (D_1, \dots, D_I)$, $\vec{d} = (d_1, \dots, d_I)$, random $\vec{X} = (X_1, \dots, X_J)$, and $\vec{x} = (x_1, \dots, x_J)$. Further, there exists values d_i^{min}, d_i^{max} , x_j^{min}, x_j^{max} , with $P(d_1^{min} \leq D_1 \leq d_1^{max}, \dots, d_I^{min} \leq D_I \leq d_I^{max}, x_1^{min} \leq X_1 \leq x_1^{max}, \dots, x_J^{min} \leq X_J \leq x_J^{max}) = 1$. We assume independence to capture the case where we cannot learn anything about upcoming emergency demand by observing elective demand.

2.3. The Hospital

To estimate the number of free units of resource r in the coming days, it is sufficient to know all states of all currently admitted and scheduled patients. Following this thought, we model the state of the hospital as a vector of states $\vec{z} = (z_1, z_2, \dots)$ of patients that were admitted to the hospital in previous periods and elective patients that are already scheduled but not yet admitted. We will refer to the k th element of \vec{z} as the health state of patient number k . Our assumption that patients stay for a maximum of N time periods in the hospital and the way we choose the indices of newly arriving patients ensure that the dimensionality of \vec{z} does not grow to infinity over time and can be restricted to a fixed finite value (as shown in Theorem 1 below).

Using this notation, we let the state of the hospital on day n be the vector of patient states after they have evolved to day n but before the emergency and elective patients have been admitted that day. As a consequence, not all patient states $z \in \mathcal{Z}'$ can actually be observed, but only those states z for which there exists an $z' \in \mathcal{Z}'$ with $p_{z'z} > 0$. Denote the set of all such states by $\mathcal{Z}_0 \subseteq \mathcal{Z}'$. So if $t_i = 0$, one might have $0_{j_i} \notin \mathcal{Z}_0$. Further let $\mathcal{Z} = \mathcal{Z}' \cup \bigcup_{r=1}^R \{1_{I+r}\}$ with $u_r(1_{I+r}) = 1$ and $u_{r'}(1_{I+r}) = 0$ for all $r' \neq r = 1, \dots, R$. (The states 1_{I+r} can be viewed as health states of artificial patients and are introduced for technical reasons, which we explain in detail in the proof of Theorem 3.)

2.4. Expected Contribution and Overbooking

If the demand by elective patients is given by \vec{d} and the hospital accepts $a_{i\tau}$ patients of type (j_i, f_i, t_i) for admission in τ days with $\tau = 0, \dots, t_i$, $i = 1, \dots, I$, it must hold that $\sum_{\tau=0}^{t_i} a_{i\tau} \leq d_i$, $a_{i\tau} \in \mathbb{N}_0$ for such an action to be feasible. If we let f_i denote the expected contribution of an admitted elective patient, the expected contribution of action $a_{i\tau}$ is hence $\sum_{i=1}^I f_i \sum_{\tau=0}^{t_i} a_{i\tau}$.

Since our extended definition of emergency patients covers teaching and research cases, elective patients can be purely viewed as a source of income. Thus, our goal is to find a dynamic admission and scheduling policy for elective patients that depends on current hospital utilization with the aim of maximizing the expected contribution minus penalty costs in the presence of emergency patients. Since the time window during which a newly arriving elective patient must be admitted to the hospital is usually relatively short, we do not consider discounting in our model.

On the current day, the combined resource requirements of all patients in the hospital in state \vec{z} is given by $\sum_{k=1}^{\infty} u_r(z_k)$ for $r = 1, \dots, R$. To this, we must add the resource requirements of all elective patients that will be admitted today, $\sum_{i=1}^I a_{i0} u_r(0_{j_i})$, as well as the resource requirements of the random number of emergency patients that will appear during the day, $\sum_{j=1}^J u_r(0_j) X_j$.

Writing $[x]^+ = \max\{0, x\}$, resource r is overbooked by

$$\left[\sum_{k=1}^{\infty} u_r(z_k) + \sum_{i=1}^I a_{i0} u_r(0_{j_i}) + \sum_{j=1}^J u_r(0_j) X_j - c_r \right]^+$$

units on a day with hospital state \vec{z} .

The expected contribution from elective patients minus overbooking costs add up to

$$\sum_{i=1}^I f_i \sum_{\tau=0}^{t_i} a_{i\tau} - \sum_{r=1}^R \pi_r E \left(\left[\sum_{k=1}^{\infty} u_r(z_k) + \sum_{i=1}^I a_{i0} u_r(0_{j_i}) + \sum_{j=1}^J u_r(0_j) X_j - c_r \right]^+ \right). \quad (1)$$

2.5. Dynamics

Using days as time periods, assume the hospital is in state \vec{z} in the current time period, demand \vec{d} is observed, and action $A = (a_{i\tau})_{i=1, \dots, I, \tau=0, \dots, t_i}$ is taken. We also assume that the hospital cannot influence the number of emergency patients \vec{X} with different admission diagnosis to be admitted throughout that time period. After this period's emergency and scheduled elective patients have been admitted, time moves forward, i.e. patient states evolve.

We assume that patient care is the number one priority and patients will always be treated such that the best possible treatment is guaranteed. In particular, we do not assume that hospital administration can choose actions to speed up patients' discharge for the sake of resource usage optimization. While this might be done in some situations, it is not ethical and the goal of this paper is not to recommend such actions. Patient states hence evolve purely randomly given the best care the hospital can provide. So a patient in state z will be in state z' in the next period with

probability $p_{zz'}$. Within the hospital state vector $\vec{z} = (z_1, z_2, \dots)$, we will refer to z_k as the patient state of patient k and say that patient number k is unused if $z_k = \diamond$.

Each day, a maximum of $\sum_{i=1}^I d_i^{max}$ requests for electives and $\sum_{j=1}^J x_j^{max}$ emergency requests arrive. The time between the request and discharge is no longer than $N + t_i$ time units for elective patients and no longer than N time units for emergencies. So there can never be more than $K = \sum_{i=1}^I d_i^{max}(t_i + N) + \sum_{j=1}^J x_j^{max} N < \infty$ patients scheduled for admission or currently admitted but not yet discharged. In other words, there will always be at least $\sum_{i=1}^I d_i^{max} + \sum_{j=1}^J x_j^{max}$ unused patient numbers within the first K patient numbers.

When new patients are admitted, we assign them to a random unused patient number between 1 and K . So let $\mathcal{K}_1 = \{k \in \{0, \dots, K\} \text{ with } z_k = \diamond\}$ be the set of “unused” indices, i.e. indices of discharged patients, and define $\kappa(\ell)$ be the (randomly chosen) ℓ th unused patient number, $\ell = 1, \dots, |\mathcal{K}_1|$. The values of $\kappa(\ell)$ are chosen sequentially, starting with $\ell = 1$ and $\kappa(\ell) = k$ with probability $1/|\mathcal{K}_\ell|$ for all $k \in \mathcal{K}_\ell$, where $\mathcal{K}_{\ell+1} = \mathcal{K}_\ell \setminus \{\kappa(\ell)\}$.

Let $\vec{z}'(\vec{z}, \vec{x}, A) = (z'_1, z'_2, \dots)$ be the state of the hospital after the admission and scheduling of the new elective and emergency patients as given by A and \vec{x} , but before patient states evolved to the following day's state. In particular, let

$$z'_k(\vec{z}, \vec{x}, A) = z_k \quad \text{for all } k : z_k \neq \diamond \quad (2)$$

$$z'_{\kappa(1 + \sum_{j'=1}^{j-1} x_{j'})}(\vec{z}, \vec{x}, A) = \dots = z'_{\kappa(\sum_{j'=1}^j x_{j'})} = 0_j \quad \text{for all } j = 1, \dots, J \quad (3)$$

$$\begin{aligned} z'_{\kappa(1 + \sum_{j'=1}^J x_{j'} + \sum_{i'=1}^i \sum_{\tau'=0}^{\tau_{i'}-1} a_{i'\tau'})}(\vec{z}, \vec{x}, A) &= \dots \\ &= z'_{\kappa(\sum_{j'=1}^J x_{j'} + \sum_{i'=1}^i \sum_{\tau'=0}^{\tau_i} a_{i'\tau'})}(\vec{z}, \vec{x}, A) = -\tau_i \quad \text{for all } i = 1, \dots, I, \\ &\tau_i = 0, \dots, t_i. \end{aligned} \quad (4)$$

In this intermediate state, after admission but before the health states of admitted patients evolve, the health states of previously admitted patients remain unchanged as shown in(2). Newly arriving emergency patients with diagnosis $j = 1, \dots, J$ are admitted in health state 0_j . Since there are x_j newly arriving emergency patients with diagnosis j , (3) assigns the first x_1 health states with unused indices to health state 0_1 , health states with unused indices $x_1 + 1$ to $x_1 + x_2$ to state 0_2 and so on. Since the number of elective patients of type i to be admitted in τ time periods is given by $a_{i\tau}$, (4) sets $a_{i\tau}$ health states to $-\tau_i$.

Assume that patient states evolve independently of each other, and that their evolution is independent of the arrival process of new patients. Then, when we start in state \vec{z} and do action A , there is a probability of $\sum_{\vec{x}} P(\vec{X} = \vec{x}) \cdot \prod_{k=1}^{\infty} p_{z'_k(\vec{z}, \vec{x}, A) z''_k}$ of observing hospital state $\vec{z}'' = (z''_1, z''_2, \dots)$ the following day. Finally, elective demand of the next day is observed before the new admission and scheduling decision must be made.

2.6. The Optimization Problem

The problem of maximizing expected contribution net overbooking costs can be modeled as a Markov decision process with a state space that is composed of the current state of the hospital and the demand from elective patients $(\vec{z}, \vec{d}) \in \mathcal{S} = \{((z_1, z_2, \dots), (d_1, \dots, d_I)) : z_k \in \mathcal{Z} \forall k = 1, \dots, \infty, d_i \in \mathbb{N}_0, d_i^{\min} \leq d_i \leq d_i^{\max} \forall i = 1, \dots, I\}$. The set of feasible actions is given by $\mathcal{A}(\vec{z}, \vec{d}) = \{A = (a_{i\tau})_{i=1, \dots, I, \tau=0, \dots, t_i} : a_{i\tau} \in \mathbb{N}_0, \sum_{\tau=0}^{t_i} a_{i\tau} \leq d_i \forall i = 1, \dots, I\}$. The expected one-stage reward is given by (1). Letting $q(\vec{x}, \vec{d}'')$ be the probability that emergency demand in the current period will be \vec{x} and that elective demand will be \vec{d}'' the following day, the transition probability from state (\vec{z}, \vec{d}) to state (\vec{z}', \vec{d}'') given action A is $\sum_{\vec{x}} q(\vec{x}, \vec{d}'') \cdot \prod_{k=1}^{\infty} p_{z'_k(\vec{z}, \vec{x}, A)z''_k}$.

Let $\{(\vec{z}^n, \vec{d}^n), A^n\}_{n=1,2,\dots}$ with actions $A^n = (a_{i\tau}^n)_{n=1,2,\dots, i=1, \dots, I, \tau=0, \dots, t_i}$ denote an infinite sequence of state-action pairs, $\{X_j^n\}_{j=1, \dots, J, n=1,2,\dots}$ an infinite sequence of emergency demand and let $\phi : \mathcal{S} \rightarrow \mathcal{A}$ be the decision function that specifies an action $A \in \mathcal{A}(\vec{z}, \vec{d})$ for every state in \mathcal{S} . Define the long-run average contribution net overbooking costs of the system under decision function ϕ , starting from initial state (\vec{z}^0, \vec{d}^0) as

$$\begin{aligned} \bar{J}(\phi, \vec{z}^0, \vec{d}^0) &= \limsup_{N' \rightarrow \infty} \frac{1}{N'} \sum_{n=1}^{N'} E \left[\sum_{i=1}^I f_i \sum_{\tau=0}^{t_i} a_{i\tau}^n \right. \\ &\quad \left. - \sum_{r=1}^R \pi_r \left[\sum_{k=1}^{\infty} u_r(z_k^n) + \sum_{i=1}^I a_{i0}^n u_r(0_{j_i}) + \sum_{j=1}^J X_j^n u_r(0_j) - c_r \right] \mid (\vec{z}^0, \vec{d}^0) \right]. \end{aligned}$$

We refer to this expression as the long-run time-average net contribution in the following. Using this notation, we can formulate the decision maker's problem to find an optimal, average net contribution maximizing decision rule ϕ^* from starting state (\vec{z}^0, \vec{d}^0) ,

$$J(\vec{z}^0, \vec{d}^0) = \sup_{\phi: \mathcal{S} \rightarrow \mathcal{A}} J(\phi, \vec{z}^0, \vec{d}^0). \quad (5)$$

Table 1 summarizes the salient notation in this paper. The following theorem will be useful in analyzing the problem in greater detail.

THEOREM 1. *The optimal average net contribution $J(\vec{z}^0, \vec{d}^0)$ is the same for all initial states, i.e. $J(\vec{z}, \vec{d}) = J^*$ for all $(\vec{z}, \vec{d}) \in \mathcal{S}$. Without loss of optimality, the action space can be reduced to actions in $\mathcal{A}^K(\vec{z}, \vec{d}) = \{A = (a_{i\tau})_{i=1, \dots, I, \tau=0, \dots, t_i} : a_{i\tau} \in \mathbb{N}_0, \sum_{\tau=0}^{t_i} a_{i\tau} \leq d_i \forall i = 1, \dots, I\}$ and the state space can be reduced to a finite state space $\mathcal{S}^K := \{((z^1, \dots, z^K), \vec{d}) \in \mathcal{Z}_0^K \times \mathbb{N}_0^I : d_i^{\min} \leq d_i \leq d_i^{\max} \forall i = 1, \dots, I\}$.*

All proofs can be found in the appendix. From Theorem 1, we can conclude that the exact patient admission problem defined on \mathcal{S} has the same average net contribution (5) as the corresponding problem with restricted state and action spaces \mathcal{S}^K and \mathcal{A}^K . Thus, we can rely on the standard theory for Markov decision processes as given in Bertsekas (2007).

3. A Simple Deterministic Upper Bound

One way of approximating the decision problem is to replace all demand and resource usage data by its expected value and to set $t_i = 0$ for all patient types $i = 1, \dots, I$. Under this assumption, the expected number of emergency patient arrives in each time period and their usage equals the expected usage in each time period. Any given day, $\sum_{j=1}^J E[u_r(Z^n)|Z^0 = 0_j]E[X_j]$ units of capacity of resource r will be used by emergency patients that arrived n time periods earlier. In total, that gives $\sum_{n=0}^N \sum_{j=1}^J E[u_r(Z^n)|Z^0 = 0_j]E[X_j]$ units of resource r used by emergency patients.

Accepting a_i elective patients of type i then uses $\sum_{i=1}^I E[u_r(Z^n)|Z^0 = 0_{j_i}]a_i$ units of capacity of resource r in n time periods. In the long run, $\sum_{n=0}^N \sum_{i=1}^I E[u_r(Z^n)|Z^0 = 0_{j_i}]a_i$ units of resource r will then be used for elective patients in each time period.

Symbol	Explanation
$A = (a_{i\tau})_{i=1, \dots, I, \tau=0, \dots, t_i}$	number of patients of type i , who asked for admission in the current period and are accepted in τ periods
c_r	daily capacity of resource r
D_i, d_i	demand from elective patients of type i and its realization
f_i	expected contribution from patient of type i
i, I	patient types $i = 1, \dots, I$
j, J	diagnosis $j = 1, \dots, J$
$J(\vec{z}, \vec{d})$	optimal average net contribution given initial state (\vec{z}, \vec{d})
n	time period
r, R	resources $r = 1, \dots, R$
π_r	per unit per day penalty for overbooking resource r
t_i	horizon during which patients of type i need to be accepted
$u_r(Z)$	number of units of resource r required by a patient in state Z
X_j, x_j	demand from emergency patients with diagnosis j and its realization
Z, z	patient health state and its realization

Table 1 Salient notation

Maximizing the expected one period net contribution gives the following decision problem with value g^{DUP} :

$$\max_{a_i, i=1, \dots, I} \sum_{i=1}^I f_i a_i - \sum_{r=1}^R \pi_r \left[\sum_{n=0}^N \left(\sum_{i=1}^I E[u_r(Z^n)|Z^0 = 0_{j_i}]a_i + \sum_{j=1}^J E[u_r(Z^n)|Z^0 = 0_j]E[X_j] \right) - c_r \right]^+ \quad (6)$$

$$s.t. \quad 0 \leq a_i \leq E[D_i] \quad \forall i = 1, \dots, I. \quad (7)$$

The objective (6) maximizes expected contribution minus penalty cost, while (7) enforces the condition that for each class of patients, we cannot accept more patients into the hospital than expected demand.

Problem (6)-(7) can easily be reformulated as a linear program and solved efficiently for large problem instances. By the following theorem, however, it only gives an upper bound on the expected net contribution. The actions recommended need not be implementable.

THEOREM 2. *Problem (6)-(7) with objective value g^{DUP} gives an upper bound to (5).*

4. Approximate Dynamic Programming and Upper Bounds

In light of Theorem 1, the patient admission problem (5) can be reduced to a problem with finite state and action space. This ensures the existence of an optimal stationary policy. It can be found by solving the average net contribution maximizing dynamic programming optimality equations (c.f. Bertsekas 2007). So let $h : \mathcal{S}^K \rightarrow \mathbb{R}$ denote the bias function and g the average net contribution. Then, the average net contribution maximizing optimality equations are

$$h(\vec{z}, \vec{d}) = \max_{A \in \mathcal{A}^K(\vec{z}, \vec{d})} \sum_{i=1}^I f_i \sum_{\tau=0}^{t_i} a_{i\tau} - \sum_{r=1}^R \pi_r E \left(\left[\sum_{k=1}^{\infty} u_r(z_k) + \sum_{i=1}^I u_r(0_{j_i}) a_{i0} + \sum_{j=1}^J u_r(0_j) X_j - c_r \right]^+ \right) - g + \sum_{z'', \vec{x}, \vec{d}''} q(\vec{x}, \vec{d}'') \prod_{k=1}^{\infty} p_{z'_k(\vec{z}, \vec{x}, A) z''_k} h(z'', \vec{d}'') \quad (8)$$

for all $(\vec{z}, \vec{d}) \in \mathcal{S}^K$ with $z'(\vec{z}, \vec{x}, A)$ given by (3)-(4).

Let g^* be the solution to (8), the maximum expected average net contribution as given in (5). It is well known that g^* is also given by the optimal solution of the following, linear programming formulation of this problem with variables g and $h : \mathcal{S}^K \rightarrow \mathbb{R}$ (Bertsekas 2007):

$$\begin{aligned} \min_{h(\vec{z}, \vec{d}), g} \quad & g \quad (9) \\ g \geq \quad & \sum_{i=1}^I f_i \sum_{\tau=0}^{t_i} a_{i\tau} - \sum_{r=1}^R \pi_r E \left(\left[\sum_{k=1}^{\infty} u_r(z_k) + \sum_{i=1}^I u_r(0_{j_i}) a_{i0} + \sum_{j=1}^J u_r(0_j) X_j - c_r \right]^+ \right) \quad (10) \\ & + \sum_{z'', \vec{x}, \vec{d}''} q(\vec{x}, \vec{d}'') \prod_{k=1}^{\infty} p_{z'_k(\vec{z}, \vec{x}, A) z''_k} h(z'', \vec{d}'') - h(\vec{z}, \vec{d}) \\ & \forall (\vec{z}, \vec{d}) \in \mathcal{S}^K, A \in \mathcal{A}^K(\vec{z}, \vec{d}). \end{aligned}$$

Problem (9)-(10), however, is difficult to solve because of its large number of variables and constraints. Consequently, we approximate the bias function h to reduce the number of variables that need to be solved.

Our approximation is based on the idea of determining the approximated marginal cost of using resource r in n time periods, denoted by V_{rn} for all $r = 1, \dots, R$ and $n = 1, \dots, N$. Those values will be helpful when comparing the resource usage and the contribution gained from different patient classes and help devise heuristics for admission and scheduling. To estimate V_{rn} for all $r = 1, \dots, R$ and $n = 1, \dots, N$, we first map the current state to the expected resource usage, and then base our bias function approximation on this vector. If the current hospital state is \vec{z} , the number of units of resource r used in the current time period (before emergency demand is accepted and elective patients are scheduled) is $\sum_{k=1}^{\infty} u_r(z_k)$. Patient states in future time periods are only

known in distribution, but the expected number of units of resource r used in n time units is $\sum_{z \in \mathcal{Z}} \sum_{k=1}^{\infty} P(Z^n = z | Z^0 = z_k) u_r(z)$, or $\sum_{k=1}^{\infty} E(u_r(Z^n) | Z^0 = z_k)$.

Weighting the approximated marginal costs of usage by the expected usage and valuing demand from class i by W_i , yields the affine function

$$h(\vec{z}, \vec{d}) \approx \sum_{i=1}^I W_i d_i - \sum_{r=1}^R \sum_{n=0}^N V_{rn} \sum_{k=1}^{\infty} E(u_r(Z^n) | Z^0 = z_k) \quad \forall (\vec{z}, \vec{d}) \in \mathcal{S}^K \quad (11)$$

with parameters $V_{rn} \geq 0$ and $W_i \in \mathbb{R}$ for all $r = 1, \dots, R$, $n = 0, \dots, N$ and $i = 1, \dots, I$, which we use to approximate the bias function $h(\vec{z}, \vec{d})$.

Employing the above approximation to the linear programming formulation can be considered as adding the additional condition (11) to the minimization problem (9)-(10). This leads to the following upper bound problem on (9)-(10) and consequently on the maximum expected average net contribution g^* :

$$(ADP) \quad \min_{g, V_{rn}, W_i, r=1, \dots, R, i=1, \dots, I, n=0, \dots, N} g \quad (12)$$

$$g \geq \sum_{i=1}^I f_i \sum_{\tau=0}^{t_i} a_{i\tau} - \sum_{r=1}^R \pi_r E \left(\left[\sum_{k=1}^{\infty} u_r(z_k) + \sum_{i=1}^I u_r(0_{j_i}) a_{i0} + \sum_{j=1}^J u_r(0_j) X_j - c_r \right]^+ \right) \quad (13)$$

$$+ \sum_{i=1}^I W_i (E(D_i) - d_i) + \sum_{r=1}^R \sum_{n=0}^N V_{rn} \left(\sum_{k: z_k \neq \diamond} E(u_r(Z^n) | Z^0 = z_k) - \sum_{k: z_k \neq \diamond} E(u_r(Z^{n+1}) | Z^0 = z_k) \right)$$

$$- \sum_{j=1}^J E[X_j] E(u_r(Z^{n+1}) | Z^0 = 0_j) - \sum_{i=1}^I \sum_{\tau=0}^{\min\{n+1, t_i\}} a_{i\tau} E(u_r(Z^{n+1-\tau}) | Z^0 = 0_{j_i})$$

$$\forall (\vec{z}, \vec{d}) \in \mathcal{S}^K, A \in \mathcal{A}^K(\vec{z}, \vec{d}).$$

We denote the solution to ADP by g^{ADP} .

Since ADP has $N \times R + I + 1$ variables and many constraints, one could try to solve it via column generation. This approach is commonly used in the ADP literature, see Adelman (2003), Adelman (2004), Adelman (2007). The subproblem can be written as a linear mixed integer problem, but the number of variables needed is large in realistic scenarios. Consequently, we suggest another relaxation in the following.

4.1. A Relaxation of the Upper Bound Problem

Consider ADP enforcing (13) for all states in \mathcal{S} instead of \mathcal{S}^K only. Since $\mathcal{S}^K \subset \mathcal{S}$, the conditions need to be satisfied for more state-action pairs and the resulting problem still gives an upper bound.

If we call this looser upper bound problem ALG, the following theorem shows that there always is an optimal solution to ALG with time-invariant values V_r .

THEOREM 3. *There exists an optimal solution to ALG with $0 \leq V_{rn} = V_r \leq \pi_r$ for all $r = 1, \dots, R$, $n \in \mathbb{N}_0$. Further, write $\vec{\alpha} = (\alpha_1, \dots, \alpha_I)$, $\vec{d} = (d_1, \dots, d_I)$, $\vec{\gamma} = (\gamma_1, \dots, \gamma_R)$, $\vec{V} = (V_1, \dots, V_R)$, and $\vec{W} = (W_1, \dots, W_I)$. Then, we can simplify ALG to*

$$\min_{g, V_r, W_i, r=1, \dots, R, i=1, \dots, I} g \quad (14)$$

$$\psi(\vec{\gamma}, \vec{\alpha}, \vec{d}, \vec{V}, \vec{W}, g) \leq 0 \quad \forall (\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X} \quad (15)$$

$$0 \leq V_r \leq \pi_r \quad \forall r = 1, \dots, R \quad (16)$$

with

$$\mathcal{X} = \{(\vec{\gamma}, \vec{\alpha}, \vec{d}) : \alpha_i, d_i, \gamma_r \in \mathbb{N}_0, 0 \leq \alpha_i \leq d_i, d_i^{min} \leq d_i \leq d_i^{max}, \gamma_r \leq c_r \forall i = 1, \dots, I, \forall r = 1, \dots, R\}$$

and

$$\begin{aligned} \psi(\vec{\gamma}, \vec{\alpha}, \vec{d}, \vec{V}, \vec{W}, g) := & \sum_{i=1}^I f_i \alpha_i - \sum_{r=1}^R \pi_r E \left(\left[\sum_{j=1}^J u_r(0_j) X_j - \gamma_r \right]^+ \right) + \sum_{i=1}^I W_i (E(D_i) - d_i) \\ & + \sum_{r=1}^R V_r \left(c_r - \gamma_r - \sum_{j=1}^J E[X_j] \sum_{n=1}^N E(u_r(Z^n) | Z^0 = 0_j) - \sum_{i=1}^I \alpha_i \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_i}) \right) - g. \end{aligned} \quad (17)$$

In the above formulation, $\alpha_i := \sum_{\tau=0}^{t_i} a_{i\tau}$ equals the total number of accepted patients of type i for given \vec{z} and A , and $\gamma_r := c_r - \sum_{k=1}^{\infty} u_r(z_k) - \sum_{i=1}^I u_r(0_{j_i}) a_{i0}$ is the number of units of resource r that are free for emergency patients in the current time period.

Then, using $x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})}$ as the dual variables of condition (15) for all $(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}$ and v_r as the dual variables of (16) for all $r = 1, \dots, R$, the dual of ALG is

$$\max_{x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})}, v_r} \sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} \left(\sum_{i=1}^I f_i \alpha_i - \sum_{r=1}^R \pi_r E \left(\left[\sum_{j=1}^J u_r(0_j) X_j - \gamma_r \right]^+ \right) \right) - \sum_{r=1}^R v_r \pi_r \quad (18)$$

$$\begin{aligned} s.t. \quad \sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} \left(c_r - \gamma_r - \sum_{j=1}^J E[X_j] \sum_{n=1}^N E(u_r(Z^n) | Z^0 = 0_j) \right. \\ \left. - \sum_{i=1}^I \alpha_i \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_i}) \right) + v_r \geq 0 \quad \forall r = 1, \dots, R \end{aligned} \quad (19)$$

$$\sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} (d_i - E[D_i]) = 0 \quad \forall i = 1, \dots, I \quad (20)$$

$$\sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} = 1 \quad (21)$$

$$x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} \geq 0 \quad \forall (\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X} \quad (22)$$

$$v_r \geq 0 \quad \forall r = 1, \dots, R. \quad (23)$$

Constraints (21) and (22) allow us to interpret $x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})}$ as the frequency of being in a state with $\vec{\gamma}$ resources reserved for emergency demand, facing demand \vec{d} of elective patients and accepting $\vec{\alpha}$. The

objective then maximizes the average contribution gained from elective patients minus the penalty costs caused by newly arriving emergency patients requiring more capacity than what was reserved for them. The last term of the objective combined with (19) ensures that additional penalty costs are incurred if the average amount of resources used by elective patients plus the average amount of resources used by already admitted emergency demand is larger than the amount that should be used for those demands, $c_r - \gamma_r$. This highlights the characterization of the corresponding primal decision variable V_r as the marginal cost of resource r . Constraint (20) corresponds to the primal variables W_i , $i = 1, \dots, I$ and ensures that average demand equals the expected value of demand.

Denote the optimal objective value of ALG by g^{ALG} . Since strong duality holds, column generation can be used to solve for g^{ALG} . This procedure starts with a small basis $\mathcal{X}' \subseteq \mathcal{X}$ that contains a feasible solution to (18)-(23). Then, (18)-(23) is solved given the basis \mathcal{X}' and the dual variables are obtained. In particular, let \vec{V} be the dual variables of (19), \vec{W} be the dual variables of (20) and g the dual variable of (21). These values are used to check if the solution fulfills primal feasibility by solving for the tightest constraint, $\max_{(\vec{\alpha}, \vec{d}, \vec{\gamma}) \in \mathcal{X}} \psi(\vec{\alpha}, \vec{d}, \vec{\gamma}, \vec{V}, \vec{W}, g)$ as defined in (17). If the solution of this subproblem is greater than 0, a constraint is violated, the values $(\vec{\alpha}, \vec{d}, \vec{\gamma})$ that obtain the maximum are added to the basis \mathcal{X}' . If the solution is less than or equal to 0, column generation stops since a primal/dual feasible and hence optimal solution was found.

Usually, an efficient solution via column generation is impossible because of the complexity of the subproblem. Note, however, that $\psi(\vec{\alpha}, \vec{d}, \vec{\gamma}, \vec{V}, \vec{W}, g)$ decomposes in r allowing us to determine the values of γ_r for all $r = 1, \dots, R$ independently of each other. Therefore, the problem can be separated into R minimization problems over γ_r and I minimization problems over the tuples (α_i, d_i) . The problems in γ_r can be viewed as newsvendor problems: For fixed r , γ_r represents a newsvendor quantity of a newsvendor that faces a demand from emergency patients $\sum_{j=1}^J u_r(0_j)X_j$, pays V_r per unit and incurs penalty cost of π_r per unit for unsatisfied demand. Thus, the understock cost is $\pi_r - V_r$ and the overstock cost are V_r . As a consequence, the optimal value of γ_r is the optimal newsvendor quantity with critical fractile $(\pi_r - V_r)/\pi_r$. The problems in (α_i, d_i) are linear. So only corner points of the feasible set can be optimal.

Denote the cumulative distribution function of $\sum_{j=1}^J u_r(0_j)X_j$ by F_r . For given \vec{W} and \vec{V} , the optimal values of the subproblem are

$$\gamma_r = \min \left\{ x \in \{0, \dots, c_r\} : F_r(x) \geq \frac{\pi_r - V_r}{\pi_r} \right\} \quad \forall r = 1, \dots, R \quad (24)$$

$$(\alpha_i, d_i) = \begin{cases} (d_i^{max}, d_i^{max}) & \text{if } f_i - \sum_{r=1}^R V_r \sum_{n=0}^N E(u_r(Z^n)|Z^0 = 0_{j_i}) > \max\{0, W_i\}, \text{ and} \\ (0, d_i^{max}) & \text{if } \max\{f_i - \sum_{r=1}^R V_r \sum_{n=0}^N E(u_r(Z^n)|Z^0 = 0_{j_i}), W_i\} < 0 \\ (d_i^{min}, d_i^{min}) & \text{if } 0 \leq f_i - \sum_{r=1}^R V_r \sum_{n=0}^N E(u_r(Z^n)|Z^0 = 0_{j_i}) < W_i \\ (0, d_i^{min}) & \text{otherwise} \end{cases} \quad \forall i = 1, \dots, I. \quad (25)$$

A column generation algorithm to solve ALG is given by the following steps:

0. Let $\mathcal{X}' = \{(\vec{\alpha}, \vec{d}, \vec{\gamma}) : \alpha_i = 0, d_i \in \{0, d_i^{max}\} \forall i = 1, \dots, I, \gamma_r = 0 \forall r = 1, \dots, R\}$ be the initial basis.
1. Solve (18)-(23) with \mathcal{X}' in place of \mathcal{X} . Let \vec{V} be the dual variables of (19), \vec{W} be the dual variables of (20) and g the dual variable of (21).
2. Determine $((\vec{\alpha}, \vec{d}, \vec{\gamma}))$ as given in (24) and (25).
3. If $\psi(\vec{\gamma}, \vec{\alpha}, \vec{d}, \vec{V}, \vec{W}, g) \leq 0$ stop. Otherwise, go back to step 1 and repeat the steps with updated basis $\mathcal{X}' = \mathcal{X}' \cup \{((\vec{\alpha}, \vec{d}, \vec{\gamma}))\}$.

The following theorem establishes the relationship between g^* and the bounds g^{ADP} , g^{ALG} and g^{DUP} .

THEOREM 4. $g^* \leq g^{ADP} \leq g^{ALG} \leq g^{DUP}$.

4.2. An Illustrative Example

To better understand the difference between the upper bounds on g^* , and to demonstrate that g^{ADP} can be a tight upper bound, we construct a stylized example where the optimal solution can be obtained analytically.

Consider the case of two-resources, $R = 2$, and three diagnosis, $J = 3$, that require only one day at the hospital each so that $p_{0_1\circ} = p_{0_2\circ} = p_{0_3\circ} = 1$. Patients with diagnosis 1 require one unit of resource 1 and patients with diagnosis 2 require one unit of resource of 2 on that day. Patients with diagnosis 3 require two units of resource 2, giving $u_1(0_1) = 1, u_2(0_1) = 0, u_1(0_2) = 0, u_2(0_2) = 1, u_1(0_3) = 0, u_2(0_3) = 2$. There are two types of elective patients, $I = 2$, and $(j_1, f_1, t_1) = (1, 3, 0)$, $(j_2, f_2, t_2) = (3, 6, 0)$. The hospital has a capacity of $c_1 = c_2 = 10$ units of the resource per day and faces a constant demand of 10 elective patients of both types per day, i.e. $P(D_1 = 10) = P(D_2 = 10) = 1$. Emergency demand arrives for diagnosis 1 and 2 independently and is equally likely to be any integer between 6 and 10 each, no emergency demand for diagnosis 3 is observed. So, $E(X_1) = E(X_2) = 8, E(X_3) = 0$. Penalty costs for not serving emergencies are $\pi_1 = \pi_2 = 12$.

Because no patient ever stays for more than one day and resource requirements between patients do not overlap, the optimal solution is given by the static one-period solution. The problem of determining the optimal number of units to reserve for emergency patients at resource 1, γ_1^* , reduces to solving a newsvendor problem. Because elective patients of type 2 do not require resource 1, the overage cost for this newsvendor equals the opportunity loss of not serving an elective patient of type 1, $f_1 = 3$. If the hospital reserves one bed less for emergencies, the hospital faces a penalty cost of 12 but gained $f_1 = 3$ from using this resource for an elective patient of type 1, so the underage cost is $12 - 3 = 9$. Thus, the critical fractile is

$$\frac{\text{underage cost}}{\text{underage cost} + \text{overage cost}} = \frac{9}{9 + 3} = 0.75.$$

Therefore, the optimal action is to reserve as much capacity as needed to serve emergency demand of each type 75% of the time. The cumulative probability distribution of demand from emergency patients is $F_1(8) = 0.6$ and $F_1(9) = 0.8$. From the newsvendor model, we know that γ_1^* is the smallest x for which the cumulative distribution is larger. Hence, $\gamma_1^* = 9$. Similarly, for resource 2, one can determine the critical fractile as $(2 \times 12 - 6)/(18 + 6) = 0.75$. Since resource 2 is demanded in pairs of two by elective patients only, it is easy to see that $\gamma_2^* = 10$ is optimal for resource 2 achieving a higher expected contribution than using $\gamma_2 = 9$ or $\gamma_2 = 8$. As a consequence, one elective patient of type 1 and zero of type 2 should be admitted in each time period, giving an average contribution of 3 and average penalty costs of $12 \times 0.2 = 2.4$, netting to 0.6. The same value is obtained when the optimality equations (8) are solved, so $g^* = 0.6$.

Solving (DUP) yields $a_1 = 2$ and $a_2 = 1$, so $g^{DUP} = 2 \times 3 + 1 \times 6 = 12$.

The optimal solution of (ADP) is $W_1^{ADP} = W_2 = 0$, $V_{1,0}^{ADP} = V_{2,0} = 3$, $g^{ADP} = 0.6$, which equals g^* . Solving g^{ALG} , we obtain $g^{ALG} = 1.2$. The difference between g^{ALG} and g^{ADP} can be explained by the fact that when ALG is solved, the constraints need to hold for a state that is in \mathcal{S} but not in \mathcal{S}^K . (A state with $\gamma_1 = \gamma_2 = 9$, can only be obtained by accepting artificial patients as introduced in Section 2.3.) In other words, by accepting patients from types $i = 1, \dots, I$ only, the hospital can never be in a state with

$$9 = \gamma_r = c_r - \sum_{k=1}^{\infty} u_r(z_k) - \sum_{i=1}^I u_r(0_{j_i}) a_{i0} = 10 - 0 - \sum_{i=1}^I 2 \times a_{i0},$$

with $a_{i0} \in \mathbb{N}_0$.

To summarize, in this example we have $g^* = g^{ADP} = 0.6 < g^{ALG} = 1.2 < g^{DUP} = 12$. A simple example with $g^* = g^{ADP} = g^{ALG} < g^{DUP}$ can be constructed by setting $u_2(0_3) = 1$.

5. Heuristics and Lower Bounds

In this section, we introduce different heuristics for patient admission, which are based on the upper bound problems developed in the previous section. These heuristics serve as lower bounds on the maximum net contribution g^* .

The first heuristic is directly based on the dynamic programming framework we introduced earlier. This type of heuristic is known as price-directed heuristic in the approximate dynamic programming literature (Adelman 2003, 2004). The second heuristic, which we call the newsvendor heuristic, uses insights gained from the lower bound problem ALG (14)-(16) and builds on ideas already used in practice.

5.1. Price-directed Heuristics

In approximate dynamic programming, the usual way to obtain a policy is to use the approximation in the optimality equation (8) and find the action that achieves its maximum. Due to the interpretation of the approximation parameters as prices, these heuristics are typically called price-directed policies. If we use superscripts *ADP* to denote the optimal solution of problem ADP, a price directed policy based on ADP would always choose a feasible action $A \in \mathcal{A}(\vec{z}, \vec{d})$ that maximizes

$$\begin{aligned} & \sum_{i=1}^I a_{i0} (f_i - \sum_{r=1}^R \sum_{n=1}^N V_{r(n-1)}^{ADP} E(u_r(Z^n) | Z^0 = 0_{j_i})) + \sum_{i=1}^I \sum_{\tau=1}^{t_i} a_{i\tau} (f_i - \sum_{r=1}^R \sum_{n=0}^N V_{r(n+\tau-1)}^{ADP} E(u_r(Z^n) | Z^0 = 0_{j_i})) \\ & - \sum_{r=1}^R \pi_r E \left(\left[\sum_{k=1}^{\infty} u_r(z_k) + \sum_{i=1}^I a_{i0} u_r(0_{j_i}) + \sum_{j=1}^J u_r(0_j) X_j - c_r \right]^+ \right). \end{aligned} \quad (26)$$

If we use the values obtained by solving ALG instead, this simplifies to

$$\begin{aligned} & \sum_{i=1}^I a_{i0} (f_i - \sum_{r=1}^R V_r^{ALG} \sum_{n=1}^N E(u_r(Z^n) | Z^0 = 0_{j_i})) + \sum_{i=1}^I \sum_{\tau=1}^{t_i} a_{i\tau} (f_i - \sum_{r=1}^R V_r^{ALG} \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_i})) \\ & - \sum_{r=1}^R \pi_r E \left(\left[\sum_{k=1}^{\infty} u_r(z_k) + \sum_{i=1}^I a_{i0} u_r(0_{j_i}) + \sum_{j=1}^J u_r(0_j) X_j - c_r \right]^+ \right). \end{aligned} \quad (27)$$

The terms $\sum_{r=1}^R V_{rn}^{ADP} \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_i})$ in (26) and $\sum_{r=1}^R V_r^{ALG} \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_i})$ in (27) can be viewed as the total approximate expected opportunity cost from accepting a patient of type i . We refer to the difference between the contribution of a type i patient and their total approximate expected opportunity cost as the approximated net contribution of patient i . A negative net contribution indicates that the contribution of this patient type is lower than the approximated opportunity costs of their resource usage. Since today's cost of capacity are already accounted for in the penalty costs, summation only starts at $n = 1$ when assessing the opportunity costs for patients that are accepted today.

Since in ADP, the values of V_{rn}^{ADP} may vary in n , (26) provides some guidance for scheduling decisions. In most problems of realistic size, even the computation of V_{rn}^{ADP} is difficult, however, and we will often only have V_r^{ALG} .

It is easy to see that (27) would never schedule patient types with negative net contribution for a future time slot. They may, however, be accepted immediately if current capacity is ample and $f_i - \sum_{r=1}^R V_r^{ALG} \sum_{n=1}^N E(u_r(Z^n) | Z^0 = 0_{j_i}) > 0$. In other words, patient types with negative net contribution may be accepted right now if the resources they need today will most likely not be used and their contribution is greater than the opportunity cost of their total resource usage in future time periods. Following (27) does not provide any guidance for scheduling beyond this accept now vs. accept later distinction and does not consider the future utilization of the hospital.

When implementing this heuristic, we will prioritize patient types in the order of decreasing net contribution and first determine how many to accept now and then accept as many as possible without overbooking for $\tau = 1, \dots, t_i$ (in expectation). No capacity is reserved for emergencies in future time periods. We will refer to this heuristic as the “price-directed” heuristic (PD) in our computational analysis.

5.2. The Newsvendor Heuristic

Although the above heuristic is simple to implement, it still requires the daily solution of a mixed integer LP to solve (26) or (27). In a practical implementation, an intuitive heuristic that does not require the repeated solution of an optimization problem may be preferable.

In our discussion of the relaxed lower bound problem ALG, we developed the following two ideas: First, if we extend the state space to \mathcal{S} , the approximated marginal opportunity cost of resource r , V_{rn} is independent of the time index n and equals V_r^{ALG} . If we refer to the patient type with the i^{th} highest net contribution as patient type (i) , patients of type (i) seem to be more valuable than patients of type $(i+1)$. Therefore, it is logical to prioritize patients based on their net contribution. (Depending on the current state of the system, this ranking might not represent the actual ranking of their value to the hospital since the approximated values V_r^{ALG} do not depend on the current system state.)

Second, given the values V_r^{ALG} we can approximate the hospital’s decision how much of resource r to reserve for emergency demand by a cost-minimizing newsvendor problem with demand coming from emergency patients. This demand can either be served by reserving units at a variable costs V_r^{ALG} or by overtime leading to penalty costs of π_r . Calculating the critical fractile as given in (24) yields γ_r , the number of units of resource r that need to be reserved for newly arriving emergency patients in each time period.

A straightforward implementation of these ideas prioritizes patients according to their net contribution and schedules them as early as possible given that γ_r units of resources $r = 1, \dots, R$ are still free and reserved for newly incoming emergency demand. This is formalized by the following heuristic, which we refer to as the “newsvendor” heuristic (NV) in our computational analysis:

Step 1. Let the current number of unused units of resource r in n time periods be c_r^n and the currently observed demand from electives \vec{d} . Let $i = 1$, $\tau = 0$ and go to Step 2.

Step 2. If $f_{(i)} - \sum_{r=1}^R V_r^{ALG} \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_{(i)}}) < 0$, stop. If $f_{(i)} - \sum_{r=1}^R V_r^{ALG} \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_{(i)}}) \geq 0$, then repeat the following two steps until a stopping criterion is reached:

a. If $d_{(i)} = 0$ or $\tau > t_{(i)}$ go to step 3. Otherwise, accept as many elective patients of type (i) at time τ as you can without using any “reserved” units of capacity, i.e. if $d_{(i)} > 0$ and $\tau \leq t_{(i)}$, let

$$a_{(i)\tau} = \max \left\{ 0, \min \left\{ d_{(i)}, \min_{r=1, \dots, R, n=1, \dots, N} \left\lfloor \frac{c_r^n - \gamma_r}{E(u_r(Z^n) | Z^0 = 0_{j_{(i)}})} \right\rfloor \right\} \right\}.$$

b. Let $c_r^n = c_r^n - a_{(i)\tau} E(u_r(Z^n) | Z^0 = 0_{j(i)})$, $d_{(i)} = d_{(i)} - a_{(i)\tau}$ and $\tau = \tau + 1$. Go to step 3.

Step 3. If $i < I$, let $i = i + 1$ and go back to step 2. If $i = I$ stop.

Observe that in this heuristic, for a given demand distribution, increasing capacity decreases the value of capacity V_r^{ALG} , which increases the critical fractile and hence increases the number of units to reserve for emergencies according to (24). The idea of reserving some capacity for emergency demand and prioritizing patients is easy to execute and to communicate as this is a natural extension of the 20% heuristic described in Section 6 that is used in practice.

6. Numerical Results

In this section, we first analyze a small example to study the quality of the bounds and the heuristics in various scenarios. This example is comparable in size to other problems studied in the admission control literature, e.g. Nadal Nunes et al. (2009). Although we can solve substantially larger problems, we first discuss such a small example since this allows the solution to all upper bound problems and a comparison with the exact solution. In the second part of this section, we apply our methods to data from the neurosurgery department of the Ronald Reagan UCLA Medical Center. In both of these examples, we solve the upper bound problems and analyze the performance of different heuristics.

From our conversation with hospital management and doctors at the UCLA Medical Center and other hospitals, we learned that two policies seem to be common in practice. In the first policy, hospital management advises that 20% of capacity should be reserved for newly arriving emergencies, the remaining capacity can be booked by other patients until it is used up. The doctors considered this policy to be ineffective as they felt that this led to too many beds being reserved for emergencies. Therefore, in the second policy, no capacity is reserved for emergencies and all capacity is used until no further elective patients can be admitted. We will refer to the latter policy as the “fill” heuristic and to the policy suggested by management as the “20%” heuristic in the following.

To benchmark our bounds and our heuristics, we compare their performance to the “fill” as well as to the “20%” heuristic, prioritizing patients according to their contribution f_i . In addition, we present the results of a greedy strategy of maximizing one-stage costs

$$\sum_{i=1}^I a_i f_i - \sum_{r=1}^R \pi_r E \left(\left[\sum_{k=1}^{\infty} u_r(z_k) + \sum_{i=1}^I a_i u_r(0_{j_i}) + \sum_{j=1}^J u_r(0_j) X_j - c_r \right]^+ \right).$$

to decide about the current period and schedule remaining demand in future time periods if capacity is available (in expectation). This greedy strategy is similar to the price directed heuristic suggested earlier in the sense that it uses the same knowledge of the system state in the current time period. However, it does not account for future time periods since the prices of the resources, V_r^{ALG} , are

set to 0. We test this “greedy” heuristic to demonstrate the value of prices V_r^{ALG} determined by the lower bound problems in the decision making process.

6.1. A Small Example

Consider a small department of a hospital with $R = 2$ constraining resources representing OR time and beds. There are $c_1 = 15$ time units of OR capacity per day and $c_2 = 8$ beds per night. Penalty costs were assessed at $\pi_1 = 50$ and $\pi_2 = 40$. There are three different diagnosis, $J = 3$. Transition graphs of the three diagnosis are given in Figure 2. There are three different types of elective patients. Elective patients of type 1 bring a contribution of 130, have diagnosis 1 and must be admitted today or rejected, so $(j_1, f_1, t_1) = (1, 130, 0)$. Elective patients of types 2 and 3 bring a contribution of 100 and 80, respectively, have diagnosis 2 and need to be admitted today or tomorrow, or be rejected, so $(j_2, f_2, t_2) = (2, 100, 1)$ and $(j_3, f_3, t_3) = (2, 80, 1)$. Emergency patients have diagnosis 3. We assume that $P(D_i = 2) = P(D_i = 0) = 0.4$, and $P(D_i = 1) = 0.2$ for elective types 1 and 2; type 3 has a certain demand of 1. For emergency patients, $P(X_3 = 0) = 0.2$, $P(X_3 = 1) = P(X_3 = 2) = 0.4$. We refer to this setup as the stochastic evolution and stochastic demand example. The example is small enough to solve to optimality with $g^* = 175.03$. Solving the upper bound problems gives 183.33, a value 4.7% higher than g^* , for all problems, ADP, and ALG. The deterministic upper bound problem DUP yields 270.00, which is 54.3% higher than the true value.

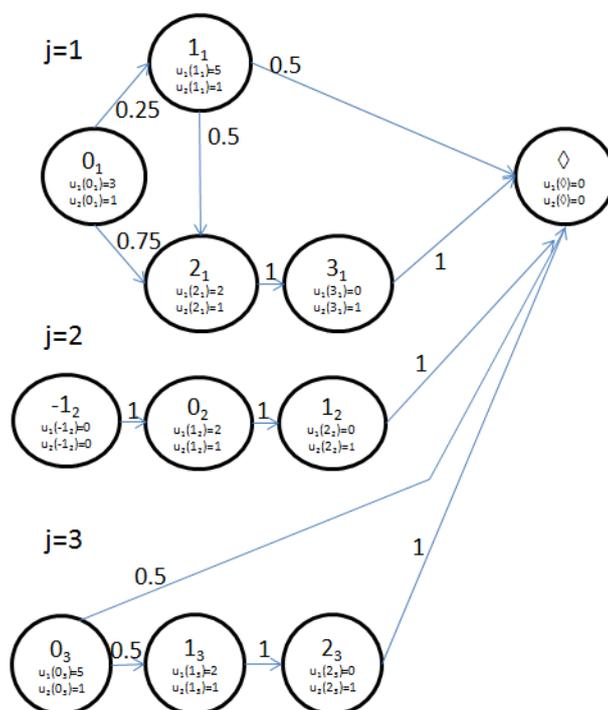


Figure 2 The Markov Chains of the three diagnosis considered in the small example.

To evaluate the admission heuristics, we simulated the admission process over 50,000 days and report the average net contribution gained. The common practice of filling the hospital up until no more patients can be accepted without using overtime yields average contribution minus penalty costs of 142, 18.9% below the optimal value g^* . The 20% heuristic performs better giving 163, or 6.8% less than the optimal policy.

Both price directed heuristics (26) and (27) yield an average value of 156, 10.7% below the optimal value g^* , but higher than a pure greedy maximization, which only gives an average net contribution of 131. The newsvendor heuristic shows an excellent performance of 173, or 1.3% below g^* .

To see how the bounds and heuristics perform across multiple scenarios, we created 19 additional scenarios based on this basic scenario with medium capacity for the OR and beds (Medium-Medium) facing a stochastic evolution of patient’s health and stochastic demand that we just introduced. First, we increased capacities by 2 each to obtain a high capacity scenario (High-High), we decreased them by 2 for the low capacity scenario (Low-Low). Combining the medium OR capacity with the low bed capacity gives the fourth scenario (Medium-Low), low OR capacity and medium bed capacity gives the fifth (Low-Medium). Using demands of 1 for all types 1, 2 and 3 and for emergencies, we create scenarios with deterministic demand. Further, we create transition graphs with the same overall resource usage but transition probabilities of 0 and 1 only to obtain scenarios with a deterministic evolution. In Tables 2 and 3 we write S in column “Ev” to mark scenarios with stochastic evolution, D for deterministic evolution. Similarly, we use S and D to determine stochastic and deterministic demand in column “De”. Solving all combinations gives a total of $2 \times 2 \times 5$ scenarios,

			High-High		Medium-Medium		Low-Low		Medium-Low		Low-Medium	
<i>Ev</i>	<i>De</i>		<i>Value</i>	$\frac{\text{Value}}{g} - 1$	<i>Value</i>	$\frac{\text{Value}}{g} - 1$	<i>Value</i>	$\frac{\text{Value}}{g} - 1$	<i>Value</i>	$\frac{\text{Value}}{g} - 1$	<i>Value</i>	$\frac{\text{Value}}{g} - 1$
S	S	g^{DUP}	310.00	43.6%	270.00	54.3%	190.00	58.6%	190.00	27.8%	245.00	77.2%
S	S	g^{ALG}	230.00	6.6%	183.33	4.7%	124.00	3.5%	164.00	10.3%	143.33	3.7%
S	S	g^{ADP}	230.00	6.6%	183.33	4.7%	124.00	3.5%	155.45	4.6%	143.33	3.7%
S	S	g^*	215.85		175.03		119.83		148.65		138.24	
S	D	g^{DUP}	310.00	52.6%	270.00	69.6%	190.00	89.7%	190.00	39.3%	245.00	103.9%
S	D	g^{ALG}	210.00	3.4%	163.33	2.6%	104.00	3.8%	144.00	5.5%	123.33	2.6%
S	D	g^{ADP}	210.00	3.4%	163.33	2.6%	104.00	3.8%	144.00	5.5%	123.33	2.6%
S	D	g^*	203.14		159.19		100.17		136.43		120.15	
D	S	g^{DUP}	310.00	3.5%	270.00	6.5%	190.00	5.6%	190.00	1.1%	245.00	6.9%
D	S	g^{ALG}	310.00	3.5%	270.00	6.5%	188.33	4.7%	190.00	1.1%	245.00	6.9%
D	S	g^{ADP}	310.00	3.5%	270.00	6.5%	188.28	4.6%	190.00	1.1%	245.00	6.9%
D	S	g^*	299.41		253.52		179.93		187.95		229.11	
D	D	g^{DUP}	310.00	0.0%	270.00	5.9%	190.00	2.7%	190.00	0.0%	245.00	9.7%
D	D	g^{ALG}	310.00	0.0%	270.00	5.9%	188.33	1.8%	190.00	0.0%	245.00	9.7%
D	D	g^{ADP}	310.00	0.0%	266.67	4.6%	185.00	0.0%	190.00	0.0%	223.33	0.0%
D	D	g^*	310.00		255.00		185.00		190.00		223.33	

Table 2 Upper bounds in different scenarios.

Ev	De	High-High		Medium-Medium		Low-Low		Medium-Low		Low-Medium		
		Value	$1 - \frac{Value}{g^*}$	Value	$1 - \frac{Value}{g^*}$	Value	$1 - \frac{Value}{g^*}$	Value	$1 - \frac{Value}{g^*}$	Value	$1 - \frac{Value}{g^*}$	
S	S	Fill	198	8.3%	142	18.9%	81	32.3%	81	32.3%	121	18.8%
S	S	20%	214	0.7%	163	6.8%	97	18.7%	97	18.7%	123	16.9%
S	S	NV	213	1.5%	173	1.3%	115	3.7%	115	3.7%	140	5.8%
S	S	PD-ALG	209	3.2%	156	10.7%	103	13.7%	103	13.7%	134	9.8%
S	S	PD-ADP	207	4.1%	156	10.7%	103	13.7%	103	13.7%	83	44.4%
S	S	Greedy	193	10.6%	131	25.0%	57	52.6%	57	52.6%	87	41.3%
S	D	Fill	180	11.3%	122	23.3%	60	39.9%	60	39.9%	102	25.3%
S	D	20%	197	2.9%	147	7.5%	87	13.3%	87	13.3%	120	11.7%
S	D	NV	197	3.1%	155	2.4%	98	2.4%	98	2.4%	129	5.1%
S	D	PD-ALG	197	2.9%	149	6.5%	92	8.6%	92	8.6%	129	5.3%
S	D	PD-ADP	199	2.1%	149	6.5%	91	9.2%	91	9.2%	129	5.2%
S	D	Greedy	186	8.5%	130	18.5%	61	39.5%	61	39.5%	93	31.5%
D	S	Fill	291	2.9%	206	18.6%	160	11.1%	160	11.1%	178	5.0%
D	S	20%	288	3.8%	225	11.3%	157	12.8%	157	12.8%	157	16.3%
D	S	NV	221	26.0%	249	1.9%	167	6.9%	168	6.9%	183	2.5%
D	S	PD-ALG	292	2.4%	240	5.3%	165	8.0%	165	8.0%	178	5.4%
D	S	PD-ADP	292	2.4%	240	5.3%	148	17.6%	148	17.6%	178	5.4%
D	S	Greedy	292	2.4%	235	7.5%	148	17.6%	148	17.6%	165	11.9%
D	D	Fill	310	0.0%	220	13.7%	167	9.9%	167	9.9%	183	3.5%
D	D	20%	310	0.0%	223	12.4%	180	2.7%	180	2.7%	180	5.3%
D	D	NV	180	41.9%	100	60.8%	165	10.8%	165	10.8%	180	5.3%
D	D	PD-ALG	310	0.0%	245	3.9%	185	0.0%	185	0.0%	180	5.3%
D	D	PD-ADP	310	0.0%	253	0.7%	163	11.7%	163	11.7%	190	0.0%
D	D	Greedy	310	0.0%	220	13.7%	160	13.5%	160	13.5%	140	26.3%

Table 3 Performance of the heuristics in different scenarios. In each scenario, the best values are highlighted in bold.

Across all scenarios, we see in Table 2 that the bounds obtained by solving ADP, and ALG are within 10% of g^* and all three tend to be close. In realistic scenarios, g^* , and for big problems even g^{ADP} are difficult, if not impossible, to compute. It is hence important to note that the value of g^{ALG} often equals g^{ADP} and g^{ALG} is always substantially lower than g^{DUP} , the only other reliable bound we would have in real settings. Even in deterministic settings, the value of g^{DUP} may be higher than the other bounds since partial acceptance is allowed in DUP, whereas in reality (and in the other lower bound problems) only integer values of patients may be accepted.

Looking at the performance of the heuristics in Table 3, three things are interesting:

First, the newsvendor heuristic outperforms all other heuristics in scenarios with medium to low demand as long as demand or patient evolution is stochastic. In particular, it outperforms the price-directed control policies in many scenarios. In purely deterministic settings, its performance is poor since it heavily relies on the probability density function of capacity used by emergency patients. In high capacity scenarios, the fact that the newsvendor heuristic can only shut off certain demand types by assigning them a negative net contribution may be counterproductive. In contrast, price-directed controls may forbid to schedule some demand types in future time periods but still accept them for the current time period if there is excess capacity.

Second, the 20% heuristic performs well in high capacity scenarios. However, this is usually not the case in practice, and we can easily construct examples where this rule would do arbitrarily bad without impacting the other heuristics. (For example, consider the case when we add capacity of 100 for each resource that is almost surely going to be consumed by incoming emergency demand. The 20% heuristic would only reserve 20% of all capacity for emergency demand, which would clearly be insufficient and costs could increase arbitrarily.)

Third, when we compare the performance of the greedy heuristic with the price-directed heuristics, the value of using the parameters obtained by ALG or ADP becomes apparent. Accounting for the cost of capacity by V_r^{ADP} or V_r^{ALG} leads to higher average net contributions than greedy one-step optimization. It is instructive to note that the price directed policy based on ADP, (26), does not perform significantly better than the one based on ALG, (27). This is important since one may not be able to solve ADP in problems of realistic size as we will see in the next section.

6.2. Real Data Example

To assess how our bounds and heuristics perform in problems of realistic size, we obtained admission data from the neurosurgery department of the Ronald Reagan UCLA Medical Center. As analyzed in Duda et al. (2013), this hospital has multiple constraining resources. Depending on the patient mix, the OR time, regular beds and ICU beds could be a bottleneck in this department, so $R = 3$. There are $c_1 = 76$ fifteen minute blocks, or time units, of OR time obtained from three operating rooms with 5 hours and one with 4 hours. Further, there are $c_2 = 24$ ICU beds and $c_3 = 48$ regular beds available for neurosurgery per day. We had 6 months of data on contribution, admission diagnosis, admission type (elective vs. emergency) and capacity usage of all constraining resources. We scaled the data on contribution for elective patients to disguise the data and estimated penalties based on interviews while ensuring that an elective patient would never be admitted if all resources they are expected to require need to be provided in overtime.

If sufficient data is available, each combination of admission diagnosis, admission type (elective, emergency, trauma, urgent), insurance, severity and expected resource usage on the day of admission can be used as predictors for resource usage to naturally obtain diagnosis groups. In our data set, however, severity was only recorded at the end of the stay and many combinations of the other parameters were observed only once. Further, resource usage in our data set was heterogeneous, zero-inflated, and highly skewed. To find diagnosis groups, we therefore first binned usage data with bins chosen to match the quartiles of the corresponding distribution (cf. MacCallum et al. 2002) and then used an ordered probit regression model (c.f. Wooldridge 2001, section 15.10) with admission diagnosis, admission type, insurance and first day resource usage as explaining variables to estimate the dependent variable corresponding to the probabilities for patients belonging to each

usage bin. Weighing the conditional expected usages of each group with the probabilities yields the expected usage for each patient for each resource. Comparing those values to the intercepts of the model, which can also be interpreted as cutoff points, we obtained the usage groups for each resource. The Cartesian product of those groups across resources gave us the final groups. Using this procedure, we grouped the 775 patients, out of which 330 were electives, into $J = 34$ diagnosis groups with $I = 15$ elective patient types. Even with the rather basic grouping, the variance of resource usage within the groups is 45% lower than the total variance for OR time, 68% lower for the regular beds and 8% lower for the ICU beds; the variance of contributions is 21% lower within groups than overall.

We found in the data that emergency patients require a substantial amount of the hospital's capacity. The empirical distribution of OR time usage on the admission day of emergency patients only is depicted in Figure 3. Observe that in certain infrequent instances, even if no elective patients were ever admitted, some overbooking is unavoidable in the OR since up to 87 units of OR time might be required by newly incoming emergency patients. In addition, emergency patients may need up to 7 ICU beds and up to 4 regular beds on the day of their admission.

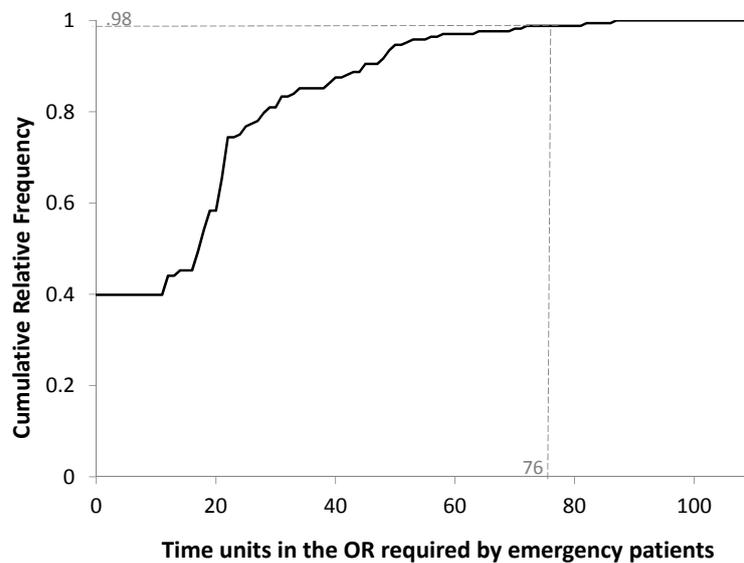


Figure 3 Empirical distribution of daily (total) OR requirements from emergency patients.

When constructing the Markov chains, we made the simplifying but realistic assumption that surgeries were performed within the first two days of the hospital stay and ICU beds were used

immediately after surgery since no data on the exact dates of resource usage was available. To estimate demand, we used the empirical distribution of emergency patients. Since only accepted elective patients can be found in hospital’s records, the exact demand from elective patients is unknown. Therefore, we generated different scenarios with differently scaled demand from elective patients to obtain an expected demand of 3 (which equals the daily average number of patients actually accepted) to 10 elective patients per time period. Since no information about the time frame of scheduling t_i , was given, we assumed that elective patients always need to be accepted within 7 days or rejected, so $t_i = 7$ for all $i = 1, \dots, I$. This gave us a total of 2281 possible health states in \mathcal{Z}_0 . The longest stay of an emergency patient was 98 days, the longest stay of an elective was 61, so we have $N = \max\{98, 61 + 7\} = 98$. Even in the lowest demand scenario with $d_i^{max} = 3$ for all $i = 1, \dots, 15$, we have $K = \sum_{i=1}^I d_i^{max}(t_i + N) + \sum_{j=1}^J x_j^{max} N > 15 \times 3 \times (7 + 98) = 4725$. Thus, the exact model (9)-(10) has more than $2281^{4725} \times 4$ variables and even more constraints. It is therefore too large to be solved directly. We solved ALG with AMPL using CPLEX 12.4 on a laptop computer with an Intel Core Duo CPU with 1.87 Ghz. This took less than one minute in all problem instances. We could not solve ADP within 60 minutes and hence only report ALG and the price directed heuristic based on the corresponding values.

Figure 4 shows the simple upper bound g^{DUP} as well as g^{ALG} for all demand scenarios. (Because of our scaling, we do not measure in \$.) Further, this figure depicts the performance of the heuristics. By “none” we refer to a setting where no elective patients would be admitted. This is an important benchmark in settings where some overbooking cannot be avoided due to highly variable emergency demand.

It is evident that both the price directed heuristic and the newsvendor heuristic perform very well and their performance is close to the upper bound. If we compute solution gaps, defined as $1 - Value/g^{ALG}$, the average solution gap of the price-directed heuristic was 5.33%, ranging from 3.72% to 10.50%, while the average solution gap for the newsvendor heuristic was 8.80%, ranging from 4.48% to 14.39%. In contrast, the average gaps for the fill, greedy, and 20% heuristic were 61.28% (between 9.09% and 119.66%), 61.77% (between 11.46% and 121.81%), and 65.61% (between 17.33% and 109.97%), respectively. It is intuitive that higher average net contribution can be achieved in higher demand scenarios. Note, however, that the performance of the greedy and practice based heuristics actually decline for higher demand scenarios due to excessive overbooking.

On further examination of the data, it was evident that for the optimized patient mix in neurosurgery, the OR time was the constraining resource. The fact that most of the uncertainty about OR requirements is resolved on the day of admission explains why the price directed and the newsvendor heuristic perform exceptionally well in this example and the gap between the upper bound g^{ALG} and their performance is small.

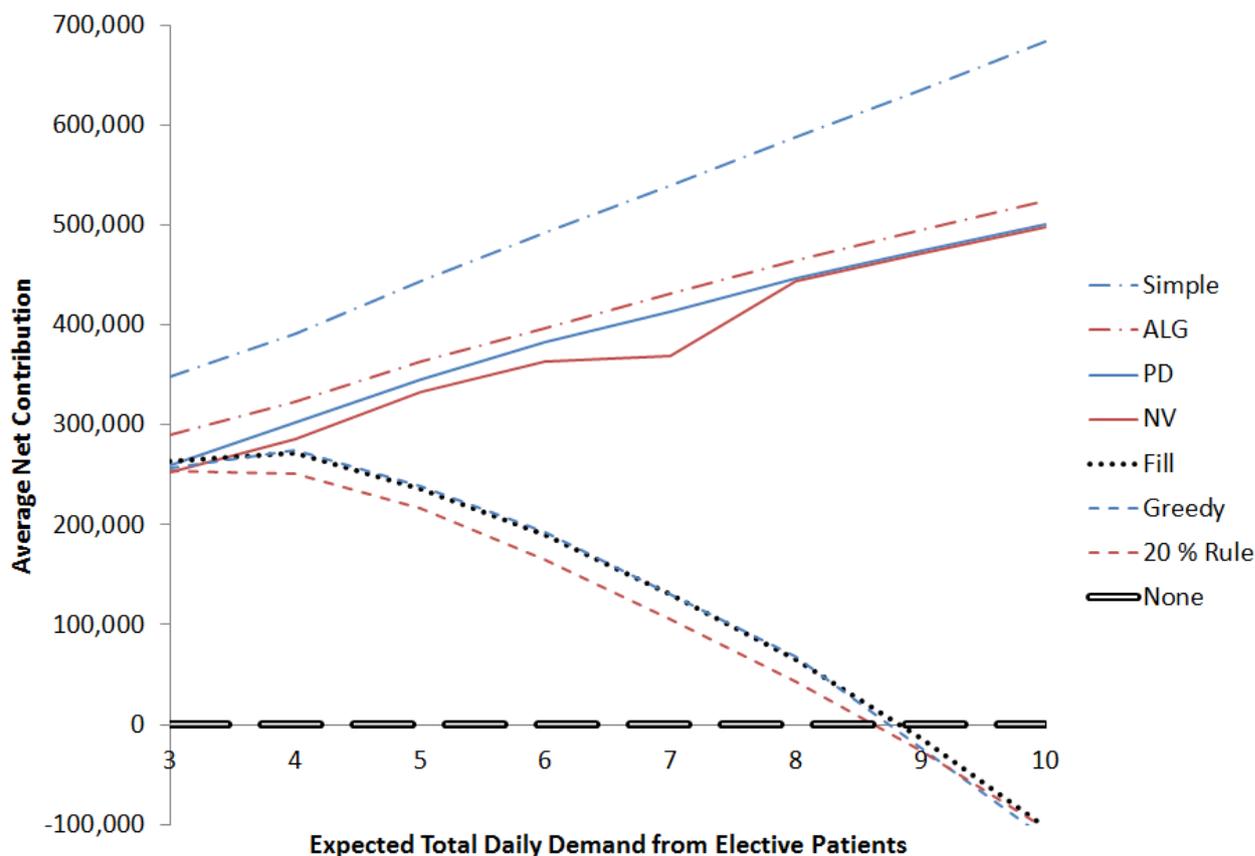


Figure 4 Bounds and performance for different demand scenarios.

Therefore, we further studied the impact of varying the number of beds on the performance of the heuristics, leaving OR capacity constant. To do this, we used the demand scaling with an expected number of 5 elective patients per day and varied c_2 and c_3 such that the expected total capacity requested over capacity available for the beds ranged from 0.6 to 1.8, in increments of 0.2. The results are depicted in Figure 5.

In all scenarios, both the price-directed heuristic PD and the newsvendor heuristic NV outperform the other heuristics by far. It is intuitive that the performance of all heuristics and the bounds decline as capacity gets smaller, i.e. the ratio gets bigger. For high ratios, the performance often yields negative values since we set the contribution from emergencies to 0 and substantial overtime is unavoidable in these cases. Again, the greedy and practice based heuristics yield very poor performance if demand vastly exceeds capacity. (The performance of greedy is almost identical to the performance of the filling heuristic; their curves are difficult to distinguish because they lie on top of each other in Figure 5.) In the scenarios with a ratio of more than one, not accepting any elective patients would leave the hospital better off than following any of those acceptance rules.

We suspect that the gap between the ALG bound and the performance of the price directed and the newsvendor heuristic widens when the capacity of the beds decline for two reasons. First, ALG replaces usage in future periods by the expected value. This was not as crucial in the real scenario, where overusage of beds hardly ever occurred due to ample capacity. As capacity gets smaller, the variability in length of stay is more important and the bound gets weaker (although it is still much stronger than the simple deterministic upper bound). Second, the performance of the heuristics decline. The price-directed heuristic directly accounts for overbooking in the current period. The cost of overbooking in future periods is only captured indirectly by exclusively accepting patients with positive net contribution. Hence, we expect its performance to decline in settings where uncertainty about overusage is resolved at a later point in time, such as in the requirements of bed capacity.

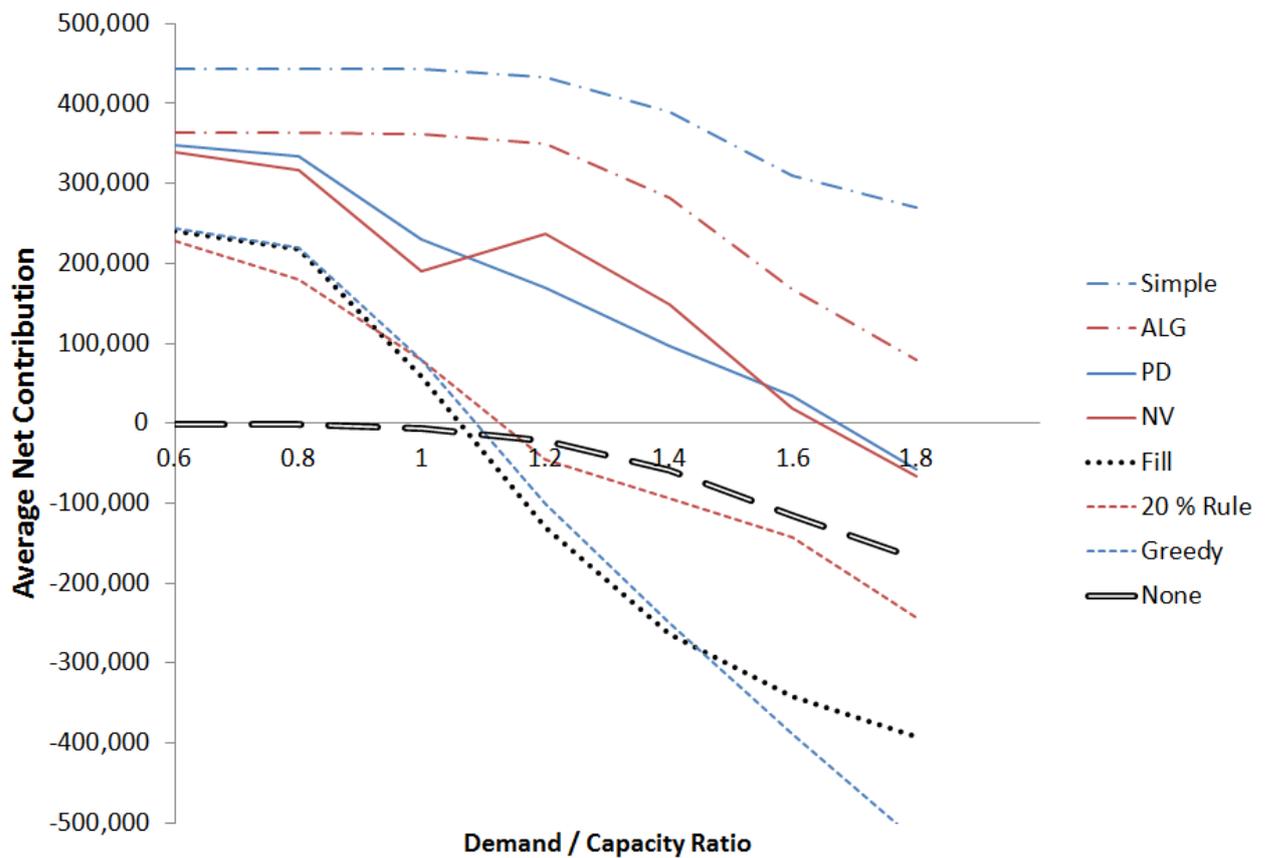


Figure 5 Bounds and performance for different bed capacities.

Across all scenarios, however, it is clear that the price-directed and the newsvendor heuristic significantly outperform other acceptance and scheduling rules.

7. Conclusion

We suggested a novel model for elective patient admission and scheduling under a stochastic evolution of patients health and care requirement with multiple resource constraints. In order to maximize expected contribution minus penalty cost, we formulated the model as a Markov decision process. Given the complexity of this model, we used techniques from approximate dynamic programming to derive an upper bound. We further simplified the upper bound problem to obtain an optimization problem that is easily solvable and yields approximated marginal values of one unit of capacity of the constraining resources. The approximated marginal values are used in heuristics for the patient admission and scheduling problem.

We find that the suggested upper bounds are significantly tighter than a naive deterministic upper bound that is obtained by replacing all stochastic elements by their expected values. Further, we show that our heuristics can improve current practice without adding capacity, by improved selection and better scheduling. The newsvendor heuristic is an extension of the “20%” heuristic, which is known among practitioners, but outperforms the latter in realistic settings. Thus, the newsvendor heuristic achieves a good balance between easy communication, intuition, and good performance.

An important takeaway for practitioners from our work is that simple rules used in practice often do not adequately account for the randomness of new arrivals, the random recovery of admitted patients, and the opportunity costs of future resource usage. Since our heuristic accounts for both the randomness of the arrival process and approximates the marginal cost of the resources, it is not surprising that it outperforms these simple heuristics, especially when uncertainty is high and capacity is tight. This increased value comes with increased data requirements. To execute the newsvendor heuristic, data is needed to estimate the demand distributions for emergency and elective patients and the expected resource usage per day for each patient type. Given the move to electronic medical data in the healthcare industry, this data should be more accessible over time.

The main contribution of our paper is to allow for multiple resource constraints and a stochastic evolution of care requirements unlike previous literature on this problem, which has mainly assumed deterministic resource requirements. While this is an important step forward, we still assume that we know the patient’s resource requirements for the current day with certainty, which might not always be the case in practice. Future research is needed to extend our results to this case. Future work could also extend our heuristics to the case of demand cancelations and time-varying arrival distributions (and resource capacities).

Acknowledgements: The authors thank the Ronald Reagan Medical Center and in particular Catherine Duda and Thomas Rosenthal for their valuable input and discussions. They further thank Neal Fultz for his assistance with the data analysis.

Appendix

Proof of Theorem 1 The fact that the model can be reduced to a model with a finite state space follows directly from the definition of K and (3)-(4) ensure that newly admitted patients or emergencies will never be assigned to indices greater than K .

So consider the reduced state space $\{(z_1, z_2, \dots), (d_1, \dots, d_I) : z_k \in \mathcal{Z}_0 \ \forall k = 1, \dots, \infty, d_i \in \mathbb{N}_0, d_i^{min} \leq d_i \leq d_i^{max} \ \forall i = 1, \dots, I\}$ and actions $\mathcal{A}^K(\vec{z}, \vec{d}) = \{A = (a_{i\tau})_{i=1, \dots, I, \tau=0, \dots, t_i} : a_{i\tau} \in \mathbb{N}_0, \sum_{\tau=0}^{t_i} a_{i\tau} \leq d_i \ \forall i = 1, \dots, I\}$ in the following.

We now show that the optimal average costs are the same for all initial states by showing that the weak accessibility condition holds, see Bertsekas (2007, p.199). The weak accessibility condition states that the state space can be partitioned into two subsets \mathcal{S}^T and \mathcal{S}^C such that 1) all states in \mathcal{S}^T are transient under every stationary policy and 2) for every two states $s', s \in \mathcal{S}^C$, there exists an integer m and a decision rule ϕ such that there is a positive probability of reaching state s' in m time periods when starting in state s ,

$$P(s^m = s' | s^0 = s, \phi) > 0.$$

In the following, we will denote by $\eta = N + \max_i t_i$ the maximum time between a patient request and their discharge.

Let \mathcal{S}^C be the set of states (\vec{z}, \vec{d}) for which there exists a sequence of η actions such that the probability of reaching (\vec{z}, \vec{d}) starting from any state $(\vec{0}, \vec{d})$ (an empty hospital) is greater than 0. All other states of the state space are $\mathcal{S}^T = \mathcal{S} \setminus \mathcal{S}^C$. Since, by definition, all patients who are scheduled for admission or admitted at time 0 will have left the hospital by time η , no other states can be visited repeatedly. So all states in \mathcal{S}^T are transient.

Now, let A^1, \dots, A^η be actions taken in states $s = s^1, s^2, \dots, s^\eta$ leading to $s^{\eta+1} = s' \in \mathcal{S}^C$ with a probability greater than 0. By the definition of \mathcal{S}^C such a sequence exists. If all states s^1, s^2, \dots, s^η are different, define $\phi(s^n) = A^n$, for all $n = 1, \dots, \eta$, choose an arbitrary feasible action for all other states. Weak accessibility follows with $m = \eta$. If one or more states are visited multiple times, pick one repeatedly occurring state at random and delete all states in between the two occurrences. Repeat this until all states in the resulting sequence $s = s''^1, s''^2, \dots, s''^M$ are different from each other. Let A''^n be the action that was chosen in state s''^n . Then, define $\phi(s''^n) = A''^n$, for all $n = 1, \dots, M$ and choose an arbitrary feasible action for all other states. Weak accessibility follows with $m = M$.

Proof of Theorem 2 The result directly follows from our proof of Theorem 4.

Proof of Theorem 3 For technical reasons, we introduce R artificial elective patient types, which represent the decision to not use one unit of resource r on the following day. These artificial types

do not bring any contribution and are numbered $i = I + 1, \dots, I + R$ with $(j_i, f_i, t_i) = (J + r, 0, 0)$. They arrive in state 0_{J+r} with $u_r(0_{J+r'}) = 0$ for all $r, r' = 1, \dots, R$ and will be in state 1_{J+r} with $u_r(1_{J+r}) = 1$ and $u_{r'}(1_{J+r}) = 0$ for all $r, r' = 1, \dots, R$ with $r \neq r'$ the following day, so that $P(Z^{n+1} = 1_{J+r} | Z^n = 0_{J+r}) = P(Z^{n+1} = \diamond | Z^n = 1_{J+r}) = 1$. Note that the introduction of such artificial patients has no impact on an optimal policy since such a policy would always reject them. To see that an optimal policy would never accept any patients from the artificial patient type, note that for any given sequence of actions A^1, A^2, \dots with $a_{I+r'}^n > 0$ for some $n, r' = 1, \dots, R$ and average net contribution J , one can easily construct a sequence with an average net contribution $J' \geq J$ by choosing A'^1, A'^2, \dots with $a_{I+r'}^n = 0$ and $a_{i\tau}^n = a_{i\tau}^n$ for all $i = 1, \dots, I, \tau = 0, \dots, t_i$ and $n \in \mathbb{N}$.

We assume that the artificial patient types $I + r$ with $r = 1, \dots, R$ have a deterministic demand of $d_{I+r}^{min} = d_{I+r}^{max} = c_r$. The model discussion and all previously proven results carry over to this setting with elective patient types $i = 1, \dots, I + R$ by extending the terms accordingly (e.g. K will now be $K = \sum_{i=1}^{I+R} d_i^{max}(t_i + N) + \sum_{j=1}^J x_j^{max} N$).

To prove that there is an optimal solution with time-invariant values V_r , we first show that $V_{rn} \leq V_{r(n-1)}$. Then we identify a subset of conditions that are the sole candidates to be the tightest constraints. Subsequently, we show how an optimal solution with time-invariant values V_r can be constructed from any given optimal solution. Finally, we show how the simplified formulation of ALG can be obtained.

First, note that we assumed that patients stay at most N time periods at the hospital. Hence, for each resource $r = 1, \dots, R$, there must be a hospital state \vec{z} for which there exists a k with $E(u_r(z_k)) > 0$ and $E(u_r(Z^n) | Z^0 = z_k) = 0$ for all $n \geq 1$. As a consequence, it is possible to choose $\sum_{k: z_k \neq \diamond} u_r(z_k)$ arbitrarily large even if $\sum_{k: z_k \neq \diamond} E(u_r(Z^n) | Z^0 = z_k) = 0$ for $n \geq 1$. For example, the hospital could be filled up with artificial patients, who all require some resources today but none tomorrow.

Further, define N_r as the maximum number of time periods into the future that resource r may be required by any patient currently admitted or scheduled.

$$N_r = \max\{n \in \{0, \dots, N\} : \exists z, z' \in \mathcal{Z} : u_r(z') > 0 \text{ and } P(Z^n = z' | Z^0 = z) > 0\}$$

and rearrange (12)-(13) to read

$$\begin{aligned} & \min_{g, V_{rn}, W_i, r=1, \dots, R, i=1, \dots, I, n \in \mathbb{N}_0} g & (28) \\ g \geq & \sum_{i=1}^I f_i \sum_{\tau=0}^{t_i} a_{i\tau} - \sum_{r=1}^R \pi_r E \left(\left[\sum_{k=1}^{\infty} u_r(z_k) + \sum_{i=1}^I u_r(0_{j_i}) a_{i0} + \sum_{j=1}^J u_r(0_j) X_j - c_r \right]^+ \right) \\ & + \sum_{i=1}^I W_i (E(D_i) - d_i) + \sum_{r=1}^R \left(V_{r0} \sum_{k=1}^{\infty} u_r(z_k) + \sum_{n=1}^{N_r} (V_{rn} - V_{r(n-1)}) \sum_{k=1}^{\infty} E(u_r(Z^n) | Z^0 = z_k) \right) \end{aligned}$$

$$\begin{aligned}
 & - \sum_{n=1}^N V_{rn} \sum_{j=1}^J E[X_j] E(u_r(Z^{n+1}) | Z^0 = 0_j) - \sum_{n=1}^N V_{rn} \sum_{i=1}^I \sum_{\tau=0}^{\min\{n+1, t_i\}} a_{i\tau} E(u_r(Z^{n+1-\tau}) | Z^0 = 0_{j_i}) \\
 & \forall (\vec{z}, \vec{d}) \in \mathcal{S}, A \in \mathcal{A}(\vec{z}, \vec{d}).
 \end{aligned} \tag{29}$$

For any solution with $V_{r0} > \pi_r$, one could always find a violated constraint by choosing $\sum_{k=1}^{\infty} u_r(z_k)$ large and $\sum_{k=1}^{\infty} E(u_r(Z^n) | Z^0 = z_k) = 0$. It follows directly that $V_{r0} \leq \pi_r$ must hold for all resources r in any feasible solution.

Since the number of patients in the hospital is unbounded, $\sum_{k=1}^{\infty} E(u_r(Z^n) | Z^0 = z_k)$ can be arbitrarily large for all $r = 1, \dots, R$, and $n = 1, \dots, N_r$. Hence, if $(V_{rn} - V_{rn-1}) > 0$, one could always find a violated constraint above. As a consequence, it must hold that

$$V_{rn} - V_{rn-1} \leq 0 \quad \forall r = 1, \dots, R, n = 1, \dots, N_r. \tag{30}$$

Now, take an optimal solution V_{rn}^* , W_i^* , and g^* with

$$0 \leq V_{rn}^* \leq V_{rn-1}^* \leq \pi_r \quad \forall r = 1, \dots, R, n = 1, \dots, N_r.$$

Since the solution is optimal, it must be feasible. Because artificial patient types require exactly one unit of a resource, there is a state $(\vec{z}, \vec{d}) \in \mathcal{S}$ with $u_r(z_k) = \lambda_r$ for any given $\lambda_r \in \mathbb{N}_0$ and $\sum_{k=1}^{\infty} E(u_r(Z^n) | Z^0 = z_k) = 0$ for all $r = 1, \dots, R$. When solving (28)-(29), it follows from (30) that the tightest constraints are given by

$$\begin{aligned}
 g^* \geq & \sum_{i=1}^I f_i \sum_{\tau=0}^{t_i} a_{i\tau} - \sum_{r=1}^R \pi_r E \left(\left[\lambda_r + \sum_{i=1}^I u_r(0_{j_i}) a_{i0} + \sum_{j=1}^J u_r(0_j) X_j - c_r \right]^+ \right) \\
 & + \sum_{i=1}^I W_i^* (E(D_i) - d_i) + \sum_{r=1}^R \left(V_{r0}^* \lambda_r \right. \\
 & \left. - \sum_{n=1}^N V_{rn}^* \sum_{j=1}^J E[X_j] E(u_r(Z^{n+1}) | Z^0 = 0_j) - \sum_{n=1}^N V_{rn}^* \sum_{i=1}^I \sum_{\tau=0}^{t_i} a_{i\tau} E(u_r(Z^{n+1-\tau}) | Z^0 = 0_{j_i}) \right).
 \end{aligned} \tag{31}$$

for any combination of demand \vec{d} , $\lambda_r \in \mathbb{N}_0$, $r = 1, \dots, R$, and actions $a_{i\tau} \in \mathbb{N}_0$ with $\sum_{\tau=0}^{t_i} a_{i\tau} \leq d_i$. Setting $g^{**} = g^*$, $W_i^{**} = W_i^*$, and choosing time-invariant values $V_{rn}^{**} = V_{r0}^*$ only reduces the value of the right hand side of the tightest constraints. If they are ensured, all other conditions are met as well, so the solution is feasible. And since $g^{**} = g^*$, the solution must be optimal, too. As a consequence, we can conclude that there always is an optimal solution with time-invariant values $0 \leq V_r^* \leq \pi_r$.

To obtain (14)-(16), use time invariant values $0 \leq V_r \leq \pi_r$ in (31), and let $\gamma_r = c_r - \lambda_r - \sum_{i=1}^I u_r(0_{j_i}) a_{i0}$ and $\alpha_i = \sum_{\tau=0}^{t_i} a_{i\tau}$ for all $r = 1, \dots, R$ and $i = 1, \dots, I$.

Proof of Theorem 4 We prove the theorem by proving three inequalities 1) $g^* \leq g^{ADP}$, 2) $g^{ADP} \leq g^{ALG}$ 3) $g^{ALG} \leq g^{DUP}$ one by one.

The first inequality, $g^* \leq g^{ADP}$, is a standard result in approximate dynamic programming and obvious if the affine approximation of the bias function (11) is understood as an additional constraint in the original linear optimization problem (9)-(10).

The second inequality $g^{ADP} \leq g^{ALG}$ follows directly from the fact that $\mathcal{S}^K \subseteq \mathcal{S}$. Since more conditions must be fulfilled the minimization problem yields a larger value.

To see the third inequality $g^{ALG} \leq g^{DUP}$, consider the dual problem (18)-(23), plug (19) into (18) and rearrange terms in (20) to obtain

$$\begin{aligned} \max_{x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})}} \sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} & \left\{ \sum_{i=1}^I f_i \alpha_i - \sum_{r=1}^R \pi_r E \left(\left[\sum_{j=1}^J u_r(0_j) X_j - \gamma_r \right]^+ \right) \right. \\ & \left. - \sum_{r=1}^R \pi_r \left[\gamma_r - c_r + \sum_{j=1}^J E[X_j] \sum_{n=1}^N E(u_r(Z^n) | Z^0 = 0_j) + \sum_{i=1}^I \alpha_i \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_i}) \right]^+ \right\} \\ \text{s.t.} \quad \sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} d_i & = E[D_i] \quad \forall i = 1, \dots, I \\ \sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} & = 1 \\ x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} & \geq 0 \quad \forall (\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}. \end{aligned}$$

Jensen's inequality yields that

$$E \left[\sum_{j=1}^J u_r(0_j) X_j - \gamma_r \right]^+ \geq \left[\sum_{j=1}^J u_r(0_j) E(X_j) - \gamma_r \right]^+.$$

Hence,

$$\begin{aligned} & \sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} \left\{ \sum_{i=1}^I f_i \alpha_i - \sum_{r=1}^R \pi_r E \left(\left[\sum_{j=1}^J u_r(0_j) X_j - \gamma_r \right]^+ \right) \right. \\ & \left. - \sum_{r=1}^R \pi_r \left[\gamma_r - c_r + \sum_{j=1}^J E[X_j] \sum_{n=1}^N E(u_r(Z^n) | Z^0 = 0_j) + \sum_{i=1}^I \alpha_i \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_i}) \right]^+ \right\} \\ & \leq \sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} \left\{ \sum_{i=1}^I f_i \alpha_i - \sum_{r=1}^R \pi_r \left(\left[\sum_{j=1}^J u_r(0_j) E(X_j) - \gamma_r \right]^+ \right. \right. \\ & \left. \left. + \left[\gamma_r - c_r + \sum_{j=1}^J E[X_j] \sum_{n=1}^N E(u_r(Z^n) | Z^0 = 0_j) + \sum_{i=1}^I \alpha_i \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_i}) \right]^+ \right) \right\} \\ & \leq \sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} \left\{ \sum_{i=1}^I f_i \alpha_i - \sum_{r=1}^R \pi_r \left(\left[\sum_{j=1}^J E[X_j] \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_j) - c_r \right. \right. \right. \\ & \left. \left. \left. + \sum_{i=1}^I \alpha_i \sum_{n=0}^N E(u_r(Z^n) | Z^0 = 0_{j_i}) \right]^+ \right) \right\} \end{aligned}$$

$$= \sum_{i=1}^I f_i a_i - \sum_{r=1}^R \pi_r \left[\sum_{n=0}^N \left(\sum_{i=1}^I E[u_r(Z^n)|Z^0 = 0_{j_i}] a_i + \sum_{j=1}^J E[u_r(Z^n)|Z^0 = 0_j] E[X_i] \right) - c_r \right]^+,$$

where the last inequality follows from $[a]^+ + [b]^+ \geq [a + b]^+$, and the equality is obtained by letting $a_i = \sum_{(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}} x_{(\vec{\gamma}, \vec{\alpha}, \vec{d})} \alpha_i$.

Further, note that for all $(\vec{\gamma}, \vec{\alpha}, \vec{d}) \in \mathcal{X}$, we have $0 \leq \alpha_i \leq d_i$. As a consequence, we can write the condition in terms of a_i as

$$0 \leq a_i = \sum_{\gamma, \alpha, d} x_{\gamma, \alpha, d} \alpha_i \leq \sum_{\gamma, \alpha, d} x_{\gamma, \alpha, d} d_i = E[D_i] \quad \forall i = 1, \dots, I.$$

So every feasible solution of (18)-(23) over \mathcal{X} can be transformed into a solution of (6)-(7) that has a larger than or equal objective value. Hence, we have $g^{ALG} \leq g^{DUP}$.

References

- Adan, I., J. Bekkers, N. Dellaert, J. Jeunet, J. Vissers. 2011. Improving operational effectiveness of tactical master plans for emergency and elective patients under stochastic demand and capacitated resources. *European Journal of Operational Research* 213(1) 290–308.
- Adan, I., J. Vissers. 2002. Patient mix optimization in hospital admission planning: a case study. *International Journal of Operations and Production Management* 22 445–461.
- Adelman, D. 2003. Price-Directed Replenishment of Subsets: Methodology and Its Application to Inventory Routing. *Manufacturing & Service Operations Management* 5(4) 348–371.
- Adelman, D. 2004. A Price-Directed Approach to Stochastic Inventory/Routing. *Operations Research* 52(4) 499–514.
- Adelman, D. 2007. Dynamic Bid Prices in Revenue Management. *Operations Research* 55 (4) 647–661.
- Adelman, D., A. Mersereau. 2008. Relaxations of weakly coupled stochastic dynamic programs. *Operations Research* 56 (3) 712–727.
- Alagoz, O., H. Hsu, A. J. Schaefer, M. S. Roberts. 2010. Markov Decision Processes: A Tool for Sequential Decision Making under Uncertainty. *Medical Decision Making* 30(4) 474–483.
- Allon, G., S. Deo, W. Lin. 2013. The impact of size and occupancy of hospital on the extent of ambulance diversion: Theory and evidence. *Operations Research* 61(3) 544–562.
- Ayvaz, N., W.T. Huh. 2010. Allocation of hospital capacity to multiple types of patients. *Journal of Revenue and Pricing Management* 9(5) 386–398.
- Bertsekas, D. P. 2005. Dynamic Programming and Suboptimal Control: A Survey from ADP to MPC. *European Journal on Control* 11(4-5) 310–334.
- Bertsekas, D. P. 2007. Dynamic programming and optimal control, Athena Scientific optimization and computation series, 2. 3rd ed. Athena Scientific.

- Bertsekas, D. P., J. N. Tsitsiklis. 1999. *Neuro-dynamic programming*, Athena Scientific optimization and computation series, 3. 2nd ed. Athena Scientific.
- Bitran, G.R., S.M. Gilbert. 1996. Managing hotel reservations with uncertain arrivals. *Operations Research* 44(1) 35–49.
- Bitran, G.R., S.V. Mondschein. 1995. An application of yield management to the hotel industry considering multiple day stays. *Operations Research* 43(3) 427–443.
- Cayirli T., E. Veral. 2003. Outpatient scheduling in health care: a review of literature. *Production and Operations Management* 12 519–549.
- Chiang, W.-C., J. Chen, X. Xu. 2007. An overview of research on revenue management: current issues and future research. *International Journal of Revenue Management* 1 97–128.
- de Farias, D. P., B. Van Roy. 2003. The linear programming approach to approximate dynamic programming. *Operations Research* 51(2) 1850–1865.
- de Farias, D. P., B. Van Roy. 2004. On constraint sampling in the linear programming approach to approximate dynamic programming. *Mathematics of Operations Research* 29(3) 462–478.
- de Farias, D. P., B. Van Roy. 2006. A cost-shaping linear program for average-cost approximate dynamic programming with performance guarantees. *Mathematics of Operations Research* 31(3) 597–620.
- Diwas, KC., C. Terwiesch. 2012. An Econometric Analysis of Patient Flows in the Cardiac Intensive Care Unit. *Manufacturing & Service Operations Management* 14(1) 50–65.
- Diwas, KC., C. Terwiesch. 2009. Impact of Workload on Service Time and Patient Safety: An Econometric Analysis of Hospital Operations. *Management Science* 55(9) 1486–1498.
- Duda, C., K. Rajaram, C. Barz, T.J. Rosenthal. 2013. A Framework on Improving Access and Customer Service Times in healthcare: Analysis and Application at the UCLA Medical Center. *Health Care Manager* 32(3) 212–226.
- Gerchak, Y., D. Gupta, M. Henig. 1996. Reservation Planning for Elective Surgery under Uncertain Demand for Emergency Surgery. *Management Science* 42(3) 321–334.
- Gocgun, Y., B. Bresnahan, A. Gbate, M. Gunn. 2011. A Markov decision process approach to multi-category patient scheduling in a diagnostic facility. *Artificial Intelligence in Medicine* 53 73–81.
- Gocgun, Y., M. Puterman. 2014. Dynamic scheduling with due dates and time windows: an application to chemotherapy patient appointment booking. *Healthcare Management Science* 17 (1) 60–76.
- Gupta, D., B. Denton. 2008. Appointment scheduling in health care: challenges and opportunities. *IIE Transactions* 40(9) 800–819.
- Helm, J.E., S. AhmadBeygi, M.P. Van Oyen. 2011. Design and analysis of hospital admission control for operational effectiveness. *Production and Operations Management* 20(3) 359–374.

- Helm, J.E., M.P. Van Oyen. 2010. Design and optimization methods for elective hospital admissions. Working Paper.
- Huh, W.T., N. Liu, V.-A. Truong. 2013. Multiresource Allocation Scheduling in Dynamic Environments, *Manufacturing & Service Operations Management* 15(2) 280–291.
- Kapadia, A.S., S.E. Vineberg, C.D. Rossi. 1985. Predicting course of treatment in a rehabilitation hospital: a Markovian model. *Computers and Operations Research* 12 459–469.
- Kim, S.H., C.W. Chan, M. Olivares G. Escobar. 2013. ICU Admission Control: An empirical study of capacity allocation and its implication on patient outcomes. Working Paper.
- Kolesar, P. 1970. A Markovian model for hospital admission scheduling. *Management Science* 16(6) 384–396.
- Kusters, R.J., P.M. Groot. 1996. Modelling resource availability in general hospitals: design and implementation of a decision support model. *European Journal of Operational Research* 88 428–445.
- Liu N., Ziya S., V. Kulkarni. 2009. Dynamic scheduling of outpatient appointments under patient no-shows and cancellations. *Manufacturing and Service Operations Management* 12 347–364
- MacCallum R.C., S. Zhang, K.J. Preacher, D.D. Rucker. 2002. On the practice of dichotomization of quantitative variables. *Psychological Methods* 7 19–40.
- Mondschein S.V., G.Y. Weintraub. 2003. Appointment policies in service operations: a critical analysis of the economic framework. *Production and Operations Management* 12 266–86.
- Naimark, D., M. D. Krahn, G. Naglie, D. A. Redelmeier and A.S. Detsky. 1997. Primer on Medical Decision Analysis: Part 5 -Working with Markov Processes. *Medical Decision Making* 17(2) 152–159.
- Nadal Nunes, L.G, S.V. de Carvalho, R. de Cassia Meneses Rodrgues. 2009. Markov decision process applied to the control of hospital elective admissions. *Artificial Intelligence in Medicine*, 47 159–171.
- OECD. 2012. Country statistical profile: United States. Country statistical profiles: Key tables from OECD. Retrieved January 4, 2014, doi: 10.1787/csp-usa-table-2011-1-en.
- Oddoye, J., D. Jones, M. Tamiz, P. Schmidt. 2009. Combining simulation and goal programming for health-care planning in a medical assessment unit. *European Journal of Operational Research* 193 250–261.
- Patrick, J., M. Puterman, M. Queyranne. 2008. Dynamic multipriority patient scheduling for a diagnostic resource. *Operations Research* 56(6) 1507–1525.
- Patrick, J. 2012. A Markov decision model for determining optimal outpatient scheduling. *Health Care Management Science* 15(2) 91–102.
- Pierskalla, W., D. Brailer. 1994. Applications of operations research in health care delivery. Pollock et al., Eds, *Handbook in OR & MS*, Springer, Heidelberg.
- Pullman, M., S. Rodgers. 2010. Capacity management for hospitality and tourism: a review of current approaches. *International Journal of Hospitality Management* 29 177–187.

- Powell, W.B. 2007. *Approximate dynamic programming: Solving the curses of dimensionality*. Wiley-Series in Probability and Statistics, Wiley, Hoboken, New Jersey.
- Robinson, L.W., R.R. Chen. 2010. A comparison of traditional and open-access policies for appointment scheduling. *Manufacturing and Service Operations Management* 12(2) 330-346.
- Saure, A., J. Patrick, S. Tyldesley, M. Puterman. 2012. Dynamic multi-appointment patient scheduling for radiation therapy. *European Journal of Operational Research* 223 573-584.
- Schaefer, A. J., M. D. Bailey, S. M. Shechter, M. S. Roberts. 2004. Modeling Medical Treatment Using Markov Decision Processes. In Brandeau, M. L., F. Sainfort und W. P. Pierskalla (eds.), *Operations Research and Health Care*. Kluwer Academic Publishers, Boston, Massachusetts, 593-612.
- Schweitzer, P. J., A. Seidmann. 1985. Generalized polynomial approximations in Markovian decision processes. *Journal of Mathematical Analysis and Applications*, 110(2), 568-582.
- Shmueli, A., C.L. Sprung, E.H. Kaplan. 2003. Optimizing admissions to an intensive care unit. *Health Care Management Science* 6 131-136.
- Social Security Advisory Board (SSAB). 2009. The unsustainable cost of health care. Retrieved January 4, 2014. http://www.ssab.gov/Documents/TheUnsustainableCostofHealthCare_graphics.pdf.
- Sonnenberg, F. A., J. R. Beck. 1993. Markov Models in Medical Decision Making: A Practical Guide. *Medical Decision Making* 13(4) 322-338.
- Sox, H. C., M. A. Blatt, M. C. Higgins, K. I. Marton. 1988. *Medical Decision Making*. Butterworths, Boston, Massachusetts.
- Talluri, K., G. van Ryzin. 2004. *The theory and practice of revenue management*. Springer, New York.
- Visser, J., I. Adan, J. Bekkers. 2005. Patient mix optimization in tactical cardiothoracic surgery planning: a case study. *IMA Journal of management Mathematics* 16(3) 281–304.
- World health organization (WHO) 2011. *World health statistics 2011*. Geneva: World Health Organization, ISBN 978-92-4-156419-9..
- Wooldridge, J.M. 2001. *Econometric analysis of cross section and panel data*, The MIT Press.